Intel® Xeon™ Processor MP

Thermal Design Guidelines

November 2002

Document Number: 298650-002

THIS DOCUMENT AND RELATED MATERIALS AND INFORMATION ARE PROVIDED "AS IS" WITH NO WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION, OR SAMPLE. INTEL ASSUMES NO RESPONSIBILITY FOR ANY ERRORS CONTAINED IN THIS DOCUMENT AND HAS NO LIABILITIES OR OBLIGATIONS FOR ANY DAMAGES ARISING FROM OR IN CONNECTION WITH THE USE OF THIS DOCUMENT.

Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications.

Intel may make changes to specifications, product descriptions, and plans at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked “reserved” or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel® Xeon™ processor MP may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling1-800-548-4725, or by visiting Intel's website at http://www.intel.com.

Intel, Pentium III, Xeon, and Intel NetBurst are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Copyright © 2002, Intel Corporation. All rights reserved.

*Other names and brands may be claimed as the property of others.

Contents

Figures

Tables

1 *Introduction*

In a system environment, the processor's temperature is a function of both the system and component thermal characteristics. The system level thermal constraints consist of the local-ambient temperature at the processor and the airflow over the processor(s) as well as the physical constraints at and above the processor(s). The processor temperature depends on the component power dissipation, size and material (effective thermal conductivity) of the integrated heat spreader (IHS); attach mechanism and the presence of a thermal cooling solution.

The continued push of technology is providing an increase in performance levels (higher operating speeds, GHz) and packaging density (more transistors). This push also increases the challenge of system thermal design. As operating frequencies increase and packaging size decreases, the power density increases and the demand on thermal cooling and system airflow increases. The result is an increased importance on system design to ensure that thermal design requirements are met for each component in the system.

The information on thermal design provided in this document is for reference only, and suggests good thermal design practices. All responsibility for determining the adequacy of any thermal or system design remains solely with the reader. Intel makes no warranties or representations that merely following all of the instructions presented in this document will result in a system with adequate thermal performance. Please refer to the terms on page ii of this document for the detailed written disclaimer.

1.1 Document Goals

The thermal power level and thermal power density of this processor generation are higher than previous Intel architecture processors. Depending on the type of system and the chassis characteristics, new system designs may be required to provide adequate cooling for the processor. The goal of this document is to provide an understanding of these thermal characteristics and discuss guidelines for meeting the thermal requirements imposed on single and multiple processor systems.

1.2 Document Scope

This document discusses thermal management and measurement techniques for the Intel® Xeon™ processor MP, which is primarily intended for server and workstation applications. The document also addresses the issues of the integrated thermal management logic and its impact on thermal design. Thermal design guidelines for the Intel® Xeon™ processor DP and Intel® Xeon™ processor with 512 KB L2 cache is covered in the *Intel® Xeon™ Processor Thermal Design Guidelines* available at http://developer.intel.com.

The physical dimensions and power numbers used in this document are for reference only. Please refer to the *Intel® Xeon™ Processor MP at 1.40 GHz, 1.50 GHz, and 1.60 GHz Datasheet* and *Intel® Xeon™ Processor MP with up to 2-MB L3 Cache on the 0.13 Micron Process Datasheet* for the product dimensions, thermal power dissipation, and maximum case temperature. In case of conflict in data, the information in the datasheet supercedes any data in this document.

1.3 References

- *Intel® Xeon™ Processor MP at 1.40 GHz, 1.50 GHz, and 1.60 GHz Datasheet*
- *Intel® Xeon™ Processor MP with up to 2-MB L3 Cache on the 0.13 Micron Process Datasheet*
- *Intel® Xeon™ Processor ProE* Models for Enabled Components*[1]
- *Intel® Xeon™ Processor IGES Models for Enabled Components*[1]
- *Intel® Xeon™ Processor Flotherm* Models*[1]
- *Intel® Xeon™ Processor MP Platform Design Guidelines*[1]
- *Intel® Xeon™ Processor and Gallatin Processor Family Thermal Test Vehicle User's Guide*
- *Intel® NetBurst™ Micro-Architecture BIOS Writers Guide*
- *603-Pin Socket Design Guidelines*[1]
- *Guidelines for Duct Design for Dual Processor Platform Applications*[2]
- *SSI Entry-Level Electronics Bay Specification*[2]
- *SSI Mid-Range Electronics Bay Specification*[2]
- *SSI High-End Electronics Bay Specification*[2]
- *European Blue Angel Recycling Standards*

1.4 Definition of Terms

- ACPI – Advanced Configuration and Power Interface (see http://www.teleport.com/~acpi).
- Bypass/no-bypass – Bypass is the area between a heatsink and any object that can act to form a duct. For this example it can be expressed as a dimension away from the outside dimension of the fins to the nearest surface.
- MSR – Model Specific Register
- $T_{LOCAL\text{-}AMBIENT}$ (T_{LA}) – The measured ambient temperature locally surrounding the processor. The ambient temperature should be measured just "upstream" of a passive heatsink, or at the fan inlet for an active heatsink.
- $T_{AMBIENT\text{-}OEM}$ – The target worst-case ambient temperature at a given external system location as defined by the system designer (OEM).
- $T_{AMBIENT\text{-}EXTERNAL}$ – The measured ambient temperature at the OEM defined external system location.
- $T_{AMBIENT\text{-}MAX}$ – The target worst-case local-ambient temperature. To determine this, place the system in a maximum external temperature environment, and measure the ambient temperature surrounding the processor. Under these conditions, $T_{LA} = T_{AMBIENT\text{-}MAX}$.
- $T_{CASE\text{-}MAX}$ – The maximum case temperature of the processor, as specified in the processor datasheet.

[1] This document is available at http://developer.intel.com.
[2] This document is available at http://www.ssiforum.org.

- T_{CASE} – The measured case temperature of the processor.
- Thermal Monitor – The Intel Xeon processor family implements a thermal management feature consisting of: an on-die thermal diode, reference current source, comparator, external bus signal, thermal control circuit and processor registers to assist with managing thermal control of the processor. Collectively, these are referred to as Thermal Monitor.
- Thermal Control Circuit – The portion of Thermal Monitor that modulates the processor's internal clocks during an over-temperature event.
- Thermal Design Power (TDP) – A processor power dissipation target derived from profiling multiple workstation and server applications. OEMs must design thermal solutions that meet or exceed the TDP as specified in the processor datasheet (also known as Thermal Design Point).
- Thermal Interface Material (TIM) – The thermally conductive compound between the heatsink and the processor case. This material fills the air gaps and voids, and enhances the spreading of the heat from the case to the heatsink.
- TBD – To be determined.
- θ_{CS} – The case to sink thermal resistance, which is dependent on the TIM. Also referred to as θ_{TIM}.
- θ_{CA} – The thermal resistance between the processor's case and the ambient air. This is defined and controlled by the system thermal solution.
- P_{MAX} – The maximum processor power, as specified in the processor's datasheet..
- 603-Pin Socket – The surface mount Zero Insertion Force (ZIF) socket designed to accept the Intel Xeon processor MP.
- U – A unit of measure used to define server rack spacing height. 1U is equal to 1.75 inches, 2U equals 3.50 inches, etc.
- Vapor Chamber – Heatsink technology similar to a heat pipe in that a liquid changes to a gaseous state to rapidly disperse the heat to a large area. A vapor chamber primary thermal benefit is to improve heat spreading over a larger surface allowing other mechanisms to remove heat from the source such as fins or pin fields.

1.5 Revision History

Revision Number	Description	Date
-002	Added Intel® Xeon™ Processor MP with up to 2-MB L3 cache information.	November 2002
-001	Initial release.	March 2002

2 Importance of Thermal Management

The objective of thermal management is to ensure that the temperature of all components in a system is maintained within functional limits. The functional temperature limit is the range within which the electrical circuits can be expected to meet their specified performance requirements. Operation outside the functional limit can degrade system performance, cause components to operate in a manner no longer guaranteed by their specified operation or cause component and/or system damage. Temperatures exceeding the maximum operating limits may result in irreversible changes in the operating characteristics of the component.

3 *Processor Packaging Technology*

The Intel Xeon processor MP uses pin grid array (PGA) package technology. Components of the package include an IHS, an organic land grid array (OLGA) or flip chip ball grid array (FC-BGA) package containing the processor die, and a pinned interposer. The IHS is designed to improve package thermal performance and is the interface for attaching a heatsink. The processor connects to the motherboard through a ZIF socket. A description of the socket can be found in the *603-Pin Socket Design Guidelines*.

The Intel Xeon processor MP utilizes a 42.5 mm (1.673 in.) OLGA core (see Figure 3-1) and the Intel Xeon processor MP with up to 2-MB L3 cache uses a 42.5 mm (1.673 in) FC-BGA core package. The processor package configurations are compatible with both the 603- and 604-pin PGA socket.

Figure 3-1. Intel® Xeon™ Processor MP Package Outline (42.5 mm OLGA or FC-BGA)

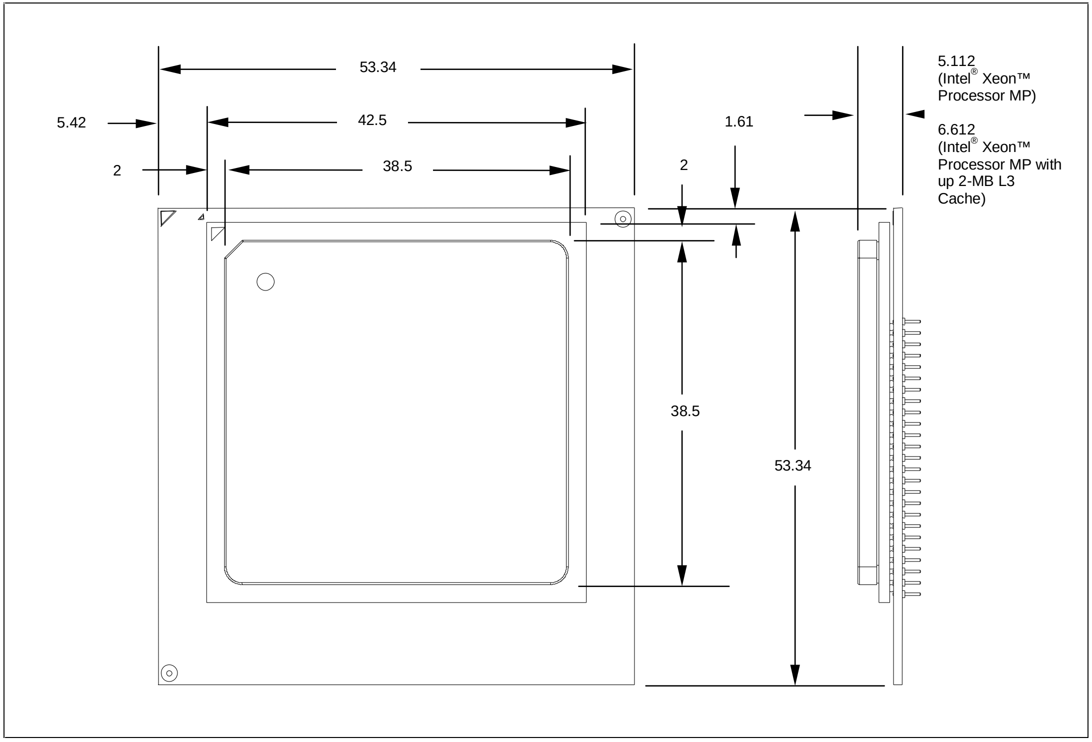

In case of conflicts in dimensions, the processor datasheet supercedes this document. All dimensions are nominal values and specified in millimeters. Please also note the difference in package height between the Intel Xeon processor MP and the Intel Xeon processor MP with up to 2-MB L3 cache. Intel implemented a 1.5 mm increase in the thickness of the IHS on the Intel Xeon processor MP with up to 2-MB L3 cache processor for enhanced thermal performance.

4 Thermal Specifications

The processor thermal power specifications can be found in the processor datasheets. Please refer to these documents for processor specifications. In order to ease the burden on chassis cooling solutions, the Thermal Monitor feature and associated logic has been integrated into the silicon of the Intel Xeon processor MP. By taking advantage of the Thermal Monitor feature, system designers may optimizethe cooling system cost and system availability while maintaining the processor reliability and performance goals. Other options within the thermal management logic allow system software to monitor and control the performance and thermal characteristics of the processor. Implementation options and recommendations are described in Section 6.

For the purposes of this application note, the following assumptions have been made about the requirements for proper operation and reliability of the processor:

- Considering the power dissipation levels and typical system local-ambient temperatures of 35°C to 45°C (95-113°F), the processor's temperatures cannot be maintained at or below its specification without additional thermal enhancement to dissipate the heat generated by the processor.
- The thermal characterization data described in later sections illustrates that both a thermal-cooling device and system airflow is needed. The size and type (passive or active) of thermal cooling device and the amount of system airflow are related and can be traded off against each other to meet specific system design constraints. In typical systems, board layout, spacing, and component placement limit the thermal solution size. Airflow is determined by the size and number of fans along with their placement in relation to the components and the airflow channels within the system. In addition, acoustic noise constraints may limit the size, number and types of fans that can be used in a particular design.

To develop a reliable, cost-effective thermal solution, all of the above variables must be considered. Thermal characterization and simulation should be carried out on the entire system, accounting for the thermal requirements of each component.

4.1 Processor Case Temperature

The IHS provides a common interface and attach location for all processor thermal solutions. The IHS improves thermal solution performance by spreading the concentrated heat from the core to a larger surface area. Thermal solutions can be active or passive. Active solutions typically incorporate a fan in the heatsink and may be smaller than a passive heatsink. Considerations in heatsink design include:

- Local-ambient temperature at the heatsink.
- Surface area of the heatsink.
- Volume of airflow over the surface area.
- Power being dissipated by the processor.
- Physical volume constraints of the system.
- Fan size, strength, and reliability (using multiple fans for redundancy).
- Air intake and exhaust.

- Altitude derating.
- System noise requirements.
- Mechanical loads and tolerances to ensure optimum thermal performance and minimum system impact.
- Fan reliability (using multiple fans for redundancy).

Techniques for measuring case temperatures are provided in Section 5.

4.2 Processor Power

The processor power, as documented in the datasheets, is the total thermal design power (TDP) that is dissipated through the IHS. This value also includes components that take into account manufacturing variations. The processor power dissipation is documented in two ways: Maximum power and TDP. Maximum power can be attained while running code specifically written to draw the most current. While running typical applications, maximum power is not usually reached, especially for a thermally significant duration of time. As a result, the TDP is provided as the thermal design target for systems. This power target is derived from profiling multiple workstation and server applications. For any excursions beyond TDP, the Thermal Monitor feature is available to maintain processor thermal specifications. Refer to Section 6 and the processor datasheets for details regarding Thermal Monitor.

5 Thermal Metrology

The following sections will discuss the techniques for testing thermal solutions. It should be noted that determining if a processor is sufficiently cooled is not as simple as it may seem. Carefully read the following instructions and determine the steps required to validate your cooling solution.

5.1 Processor Thermal Metrology

5.1.1 Thermal Resistance

The case to local-ambient (θ_{CA}) thermal resistance for a PGA package with an IHS is a measure of the cooling solution thermal performance. The case to local-ambient thermal resistance, θ_{CA}, is comprised of the case to sink (θ_{CS}) thermal resistance and the sink to local-ambient thermal resistance (θ_{SA}). Thermal resistance is measured in units of °C/W.

The case to local-ambient (θ_{CA}) thermal resistance is measured between the top of the IHS case and the local-ambient air. It is strongly dependent on the thermal conductivity and thickness of the TIM between the heatsink and surface of the processor IHS. The term θ_{SA} is a measure of the thermal resistance from the bottom of the heatsink base to the local-ambient air. θ_{SA} is dependent on the heatsink material, thermal conductivity, geometry, and strongly dependent on the air velocity through the fins of the heatsink (see Figure 5-1). θ_{CS} is the thermal resistance between the top of the IHS case and the bottom of the heatsink base. It is dependent on the TIM thermal conductivity bond-line thickness, compressive load, and the flatness and tilt of the heatsink and IHS mating surfaces.

Figure 5-1. Processor Thermal Resistance Relationships

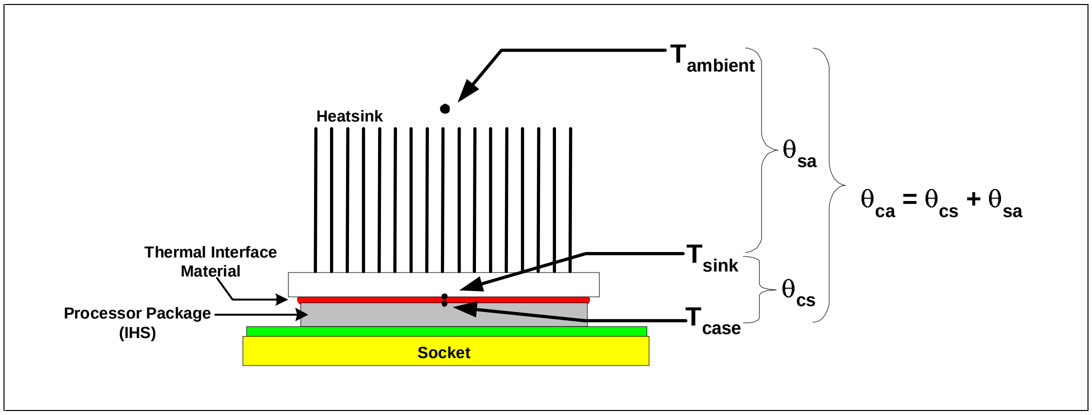

The thermal parameters are related by the following equations:

Equation 1. Thermal Case to Ambient Thermal Resistance

$$\theta_{CA} = (T_{CASE} - T_{LA}) / P_D$$

Equation 2. Thermal Case to Ambient Thermal Resistance

$$\theta_{CA} = \theta_{CS} + \theta_{SA}$$

Where:

θ_{CA}	=	case to local-ambient thermal resistance (°C/W)
T_{CASE}	=	processor case temperature (°C)
T_{LA}	=	local-ambient temperature in chassis around processor (°C)
P_D	=	processor TDP, assuming all power is dissipated through the case (W)
θ_{CS}	=	case to sink thermal resistance, dependent on the TIM (°C/W)
θ_{SA}	=	sink to local-ambient thermal resistance (°C/W)

5.1.2 Thermal Solution Performance

All processor thermal solutions attach to the processor at the IHS. The system thermal solution must adequately control the local-ambient air around the processor (T_{LA}). The lower the thermal resistance between the processor and the local-ambient air, the more efficient the thermal solution. The required θ_{CA} is dependent upon the maximum allowed processor IHS, or case temperature (T_{CASE}), the local-ambient temperature (T_{LA}), and the processor power (P_D).

Use Equation 1 and Equation 2 to determine a target θ_{CA} and θ_{SA} using the following assumptions:

T_{CASE}	=	75°C, hypothetical maximum case temperature specification
T_{LA}	=	Assume 45°C, a typical value for desktop systems
P_D	=	Assume 70 W, hypothetical TDP
θ_{CS}	=	Assume 0.12°C/W

Solving for Equation 1:

$$\theta_{CA} = (T_{CASE} - T_{LA}) / P_D$$
$$= (75 - 45) / 70$$
$$= 0.42°C/W$$

Solving for Equation 2:

$$\theta_{CA} = \theta cs + \theta_{SA}$$
$$\theta_{SA} = \theta_{CA} - \theta_{CS}$$
$$= 0.42 - 0.12$$
$$= 0.30°C/W$$

In practice, the evaluation of various thermal solutions are best evaluated at a θ_{CA} resolution level. The resulting contributions of the TIM and heatsink can be deduced after the evaluations. Given the dependency on other mechanical and thermal parameters, evaluations should also include a definition of the configuration under which the study was conducted.

5.1.3 Local-Ambient Temperature Measurement Guidelines

Local-ambient temperature, T_{LA}, is the temperature of the ambient air surrounding the processor. In a system environment, ambient temperature is the temperature of the air upstream of the processor and in its close vicinity; in an active cooling system, it is the inlet air to the active cooling device. It is necessary to determine the local-ambient temperature in the chassis around the processor to determine the thermal performance of a given thermal solution (θ_{CA}).

The local-ambient temperature is best measured as an average of the localized air surrounding the processor. The following guidelines are meant to enable accurate determination of the localized air temperature around the processor during system thermal testing.

During system thermal testing, a minimum of two thermocouples should be placed approximately 12.7 to 25.4 mm (0.5 to 1.0 in.) away from processor and heatsink as shown in Figure 5-2.

1. This placement guideline is meant to minimize localized hot spots due to the processor, heatsink, or other system components.
2. The thermocouples should be placed approximately 50.8 mm (2.0 in.) above the baseboard and 25.4 to 50.8 mm (1.0 to 2.0 in.) apart. This placement guideline is meant to minimize localized hot spots from baseboard components.
3. T_{LA} should be the average of the thermocouple measurements during system thermal testing.

Figure 5-2. Locations for Measuring Local-Ambient Temperature

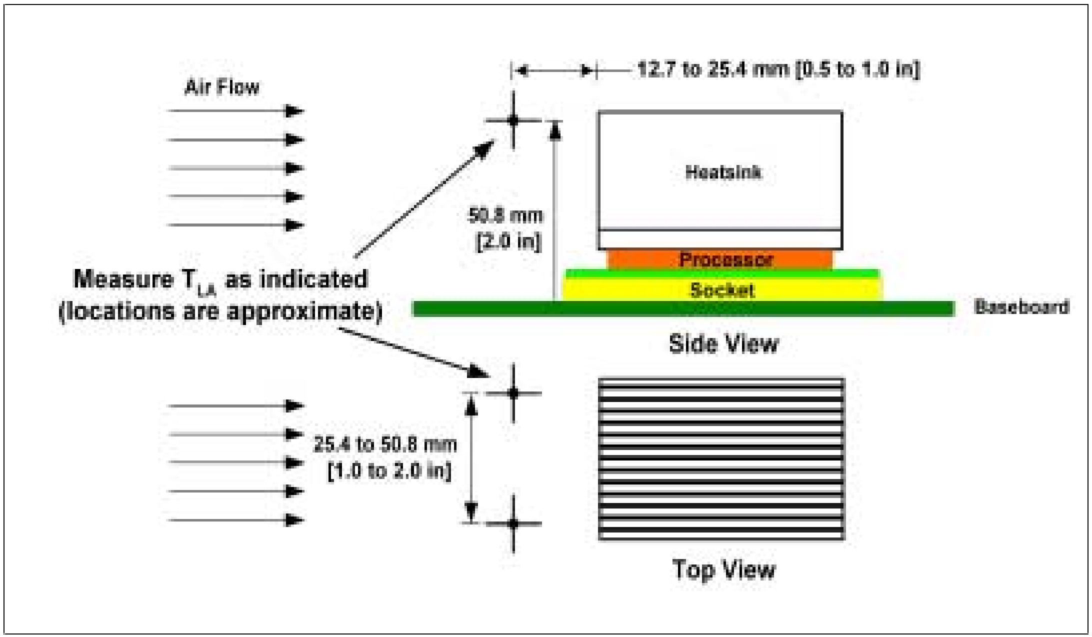

NOTE: Drawing is not to scale.

5.1.3.1 Measurements for Processor Thermal Specifications

The system integrator must make processor T_{case} measurements to determine whether a system or component thermal solution is adequate for maintaining a processor within thermal specifications. Guidelines have been established for proper techniques for measuring T_{case} temperatures. The following sections describe these guidelines for temperature measurement.

5.1.3.2 Processor Case Temperature Measurements

To ensure functionality and reliability, the processor is specified for proper operation when T_{CASE} is maintained at or below the value listed in the processor datasheets. The measurement location for T_{CASE} is the geometric center of the IHS. Figure 5-3 shows the location for T_{CASE} measurement for the Intel Xeon processor MP.

Special care is required when measuring the T_{CASE} to ensure an accurate temperature measurement. Thermocouples are often used to measure T_{CASE}. Before any temperature measurements are made, the thermocouples must be calibrated. When measuring the temperature of a surface, which is at a different temperature from the surrounding local-ambient air, errors could be introduced in the measurements. The measurement errors can be due to a poor thermal contact between the thermocouple junction and the surface of the IHS. Errors can also occur via heat loss by conduction through thermocouple leads, or by contact between the thermocouple cement and the heatsink base. To minimize these measurement errors, the following approaches are recommended.

5.1.4 Methodology for Solid Base Heatsinks (90° Angle Attach)

1. Prepare a 36 gauge or finer diameter K, T, or J type thermocouple.
2. Ensure that the thermocouple has been properly calibrated.
3. The thermocouple should be attached at a 90° angle to the IHS at the location specified for T_{CASE} measurement (see Figure 5-3 and Figure 5-4).
4. Drill a 3.8 mm (0.150 in.) maximum diameter hole through the heatsink base. This hole must be positioned on the heatsink base so that it matches with the center of the IHS when assembled. This hole will reduce the heatsink performance by 0.02°C/W.
5. Create a small depression, approximately 1.59 mm (1/16 in.) in diameter by 0.4 mm (1/64 in.) deep at the center of the IHS. This will facilitate the attach procedure by keeping the thermocouple centered and hosting the adhesive.
6. Attach the thermocouple bead or junction to the top surface of the IHS at the location specified in Figure 5-3, using high thermal conductivity cements.
7. Route the thermocouple wires through the hole in the heatsink base and attach it to the processor IHS. The use of more viscous adhesives and minimizing the use of drying accelerators will prevent problems with the adhesive spreading.
8. A small fixture may be required to hold the thermocouple and apply a steady force during the curing process to ensure the thermocouple is making contact with the IHS. A Digital Multi-Meter can be used to check continuity between the IHS and the connector as the adhesive cures.
9. Make sure there is no contact between the thermocouple cement and heatsink base. Contact will affect the thermocouple reading.
10. Verify the cured adhesive bead is smaller than 3.8 mm (0.150 in.) in diameter and height so as to fit in the hole drilled in the heatsink base. Trim as necessary.
11. Place the TIM on the heatsink base. If it is a semi-liquid type, apply it on the IHS, around the thermocouple. The clamping force will spread the TIM. If the TIM is a solid type, punch a 3.8 mm (0.150 in.) diameter hole in the center of the TIM pad and cut a line from hole to the edge corresponding to the slot for the thermocouple wire. This will allow the installation of the TIM to the IHS with the thermocouple already attached to the processor.

Figure 5-3. Processor IHS Temperature Measurement Location

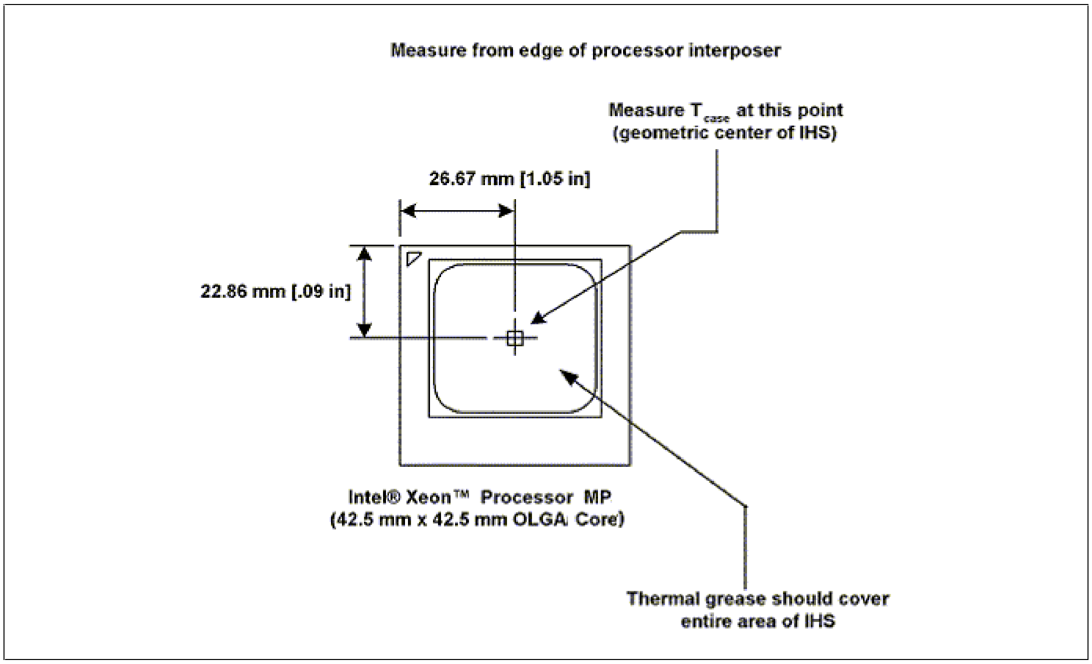

NOTE: Drawing is not to scale. In case of conflicts in dimensions, the processor datasheet supercedes this document. All dimensions are nominal values.

Figure 5-4. Technique for Measuring with 90° Angle Attachment

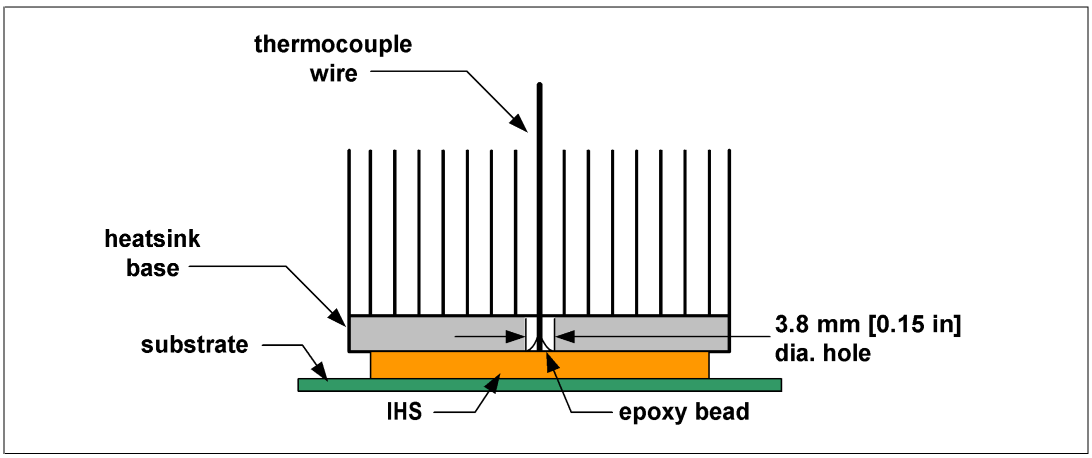

NOTE: Drawing is not to scale.

5.1.4.1 Methodology for 0° Angle Thermocouple Attach

1. Prepare a 36 gauge or finer diameter K, T, or J type thermocouple.
2. Ensure that the thermocouple has been properly calibrated.
3. Use a scribe to mark the geometric center on the topside of the IHS. This represents the location where the bead of the thermocouple will be placed. The center of the IHS can be obtained by measurement, or by drawing two diagonal lines across the length of the IHS. The cross-point will be the geometric center of the IHS. Figure 5-3 shows the thermocouple location relative to the IHS and processor interposer substrate.

4. After the marks are scribed, clean the desired thermocouple attach location with a mild solvent and a lint-free wipe or cloth. Alcohol or acetone should suffice. Remember that the cleaner the part is, the stronger the bond will be after curing.
5. Straighten the thermocouple wire by hand so that the first 100 to 150 mm (4 to 6 in.) section is reasonably straight. Use fine point tweezers to make sure that the bead and the two wires coming out are straight and untwisted. Make sure that the second layer of thermocouple insulation, sometimes clear, is not covering the bead.
6. Create a slight downward bend in the wires about 1.6 mm (1/16 in.) from the bead. Once the thermocouple is in place, this will guarantee that the thermocouple bead is making contact with the surface.
7. Place the thermocouple bead on the geometric center of the IHS (previously scribed). Apply Kapton* tape across the wire about 1/4 inch back from the bead to hold the thermocouple in place. Apply pressure to the tape to ensure a good bond. Apply additional tape pieces along the length of the wire to ensure a good temporary bond to the part. Check the electrical continuity between the thermocouple and the IHS using a multi-meter. If there is no electrical continuity between the thermocouple and the IHS, repeat steps 5 through 7.
8. With the thermocouple in place, the epoxy can now be mixed and applied. Follow the manufacturer's directions for mixing the epoxy.
9. Use a clean, finely pointed applicator to apply the epoxy to the bead. Dab the epoxy bond on the bead and the exposed wires. Use only the appropriate amount of glue to cement the thermocouple down. The entire bead should be submerged and it is best to have insulated wires protruding from the glue.
10. Additional beads of epoxy can be added, off of the IHS surface, along the length of wire to provide strain relief for the thermocouple wire. ***Only one epoxy bead should be on the IHS surface***.
11. Cure the epoxy according to the manufacturer's instructions. The cure temperature should remain below 150°F (65°C) if possible. Make sure the vibration in the oven is minimal to prevent the thermocouple bead from moving. Another alternative is to cure the epoxy at room temperature for a duration recommended by the manufacturer.
12. Once the epoxy has cured, remove all tape and check for any residual epoxy outside the thermocouple attach area on the IHS. Run the tip of your finger around the IHS surface to find any small glue dots. Remove any residual glue to prevent any impact on bond line or heatsink attach.
13. Verify the cured adhesive bead at the IHS center is smaller than 3.8 mm (0.15 in.) in diameter and 0.63 mm (0.025 in.) in height so as to fit in the groove machined in the heatsink base. Trim as necessary.
14. Check the electrical continuity between the thermocouple and the IHS again. If there is no electrical continuity between the thermocouple and the IHS, repeat steps 4 through 12.
15. Place the TIM on the heatsink base. If it is a paste, apply it on the IHS around the thermocouple. The clamping force will spread the TIM. If the TIM is a solid pad, punch a 3.8 mm (0.15 in.) diameter hole in the center of the TIM pad and cut a line from a side to the hole. This will allow the installation of the TIM to the IHS with the thermocouple already attached to the IHS.
16. In order to measure the case temperature as accurately as possibly, the heatsink must be grooved to allow room for the thermocouple wires and attach point. Depending on the heatsink, the dimensions of the groove location may vary. The system integrator should perform the appropriate analysis to define the placement dimensioning for a specific thermal design. It is imperative that the heatsink groove is aligned with the thermocouple wires and

bead. Any discrepancy will cause the heatsink to sit improperly on the IHS surface and provide erroneous data. A 1.0 mm (0.040 in.) wide groove with a depth of 0.6 mm (0.025 in.) should be milled into the heatsink base. A circular area should be milled out to accommodate the epoxy surrounding the thermocouple bead (~4.5 mm [0.18 in.] diameter, 0.6 mm [0.025 in.] deep). The center of the circular area should be located 1.3 mm (0.05 in.) off center from the location corresponding to the thermocouple bead. The offset ensures that the circular area accommodates the entire epoxy bead that covers both the thermocouple bead and the exposed thermocouple wires. See Figure 5-5 for an example of a heatsink base grooved for thermocouple installation.

Figure 5-5. Example Groove in Heatsink Base for Thermocouple Installation

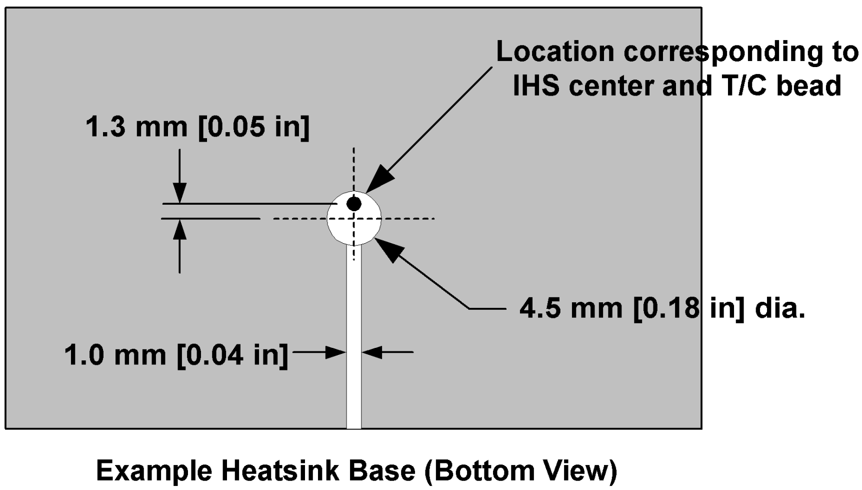

NOTE: Drawing is not to scale and is for reference only.

5.1.5 Thermal Testing Software

The Intel Xeon processor MP thermal testing software is a Win32 application that runs within a command prompt window. This software is intended for thermal evaluation purposes only and is not a general purpose application. The software does not generate the maximum processor power as defined in the processor datasheet. Details regarding the execution of the thermal testing software are provided in the “readme” file included in the software package. Contact your Intel Field Sales representative for the latest copy of the Intel® Xeon™ MP Thermal Testing Software.

6 Thermal Management Logic and New Thermal Monitor Feature

6.1 Processor Power Dissipation

An increase in processor operating frequency not only increases system performance, but also increases the processor power dissipation. The relationship between frequency and power is generalized in the following equation: $P = CV^2F$ (where P = power, C = capacitance, V = voltage, F = frequency). From this equation it is evident that power increases linearly with frequency and with the square of voltage. In the absence of power saving technologies, ever increasing frequencies will result in processors with power dissipations in the hundreds of watts. Fortunately, there are numerous ways to reduce the power consumption of a processor. Decreasing the voltage and transistor size are two examples, a third is clock modulation, which is used extensively in laptop designs.

Clock modulation is defined as periodically removing the clock signal from the processor core, which effectively reduces its power consumption to a few watts. A zero watt power dissipation level is not achievable due to transistor leakage current and the need to keep a few areas of the processor active (cache coherency circuitry, phase lock loops, interrupt recognition, etc.). Therefore, by cycling the clocks on and off at a 50% duty cycle, for example, the average power dissipation can drop by up to 50%.

Note: The processor performance will also drop by about 50% during this period, since program execution halts while the clocks are removed. Varying the duty cycle will have a corresponding influence on power dissipation and processor performance.

Laptop systems use clock modulation to control system and processor temperatures. By using various external measurement devices, laptops monitor the processor case temperature and turn on fans or initiate clock modulation to reduce processor power dissipation and ensure that all elements of the system operate within their temperature specification. Unfortunately, using external thermocouples connected to the processor package to monitor and control a thermal management solution has some inherent disadvantages. Thermal resistance (θ_{JC}) in the processor package creates a temperature delta between the processor case and silicon. This delta may be large, with the silicon temperature always being higher than the case temperature (under normal operating circumstances). When the thermocouple is measuring case temperature, not silicon temperature, significant added margin would be necessary to ensure the processor silicon does not exceed its maximum specification. Or, more clearly, clock modulation may have to be turned on when the case temperature is well below the maximum specification to ensure that the processor does not overheat. This added margin will have a substantial and unacceptable impact on system performance.

The thermal ramp rate, or change in die temperature over a specified time period ($\Delta T/\Delta t$), may be extremely high in high power processors. Ramp rates in excess of 50°C/s may occur in the course of normal operation. With this type of thermal characteristic, it would not be possible to control fans or other cooling devices based on processor core temperature. By the time the fans have spun up to speed, the processor may be well beyond a safe operating temperature. Just as large added margins would be necessary to account for package thermal gradients, equally large margins would also be necessary if temperature-controlled fans were implemented.

Clearly, a new thermal management approach is needed to support the continued increases in processor frequency and power consumption.

A new processor on-die thermal management feature resolves these issues so that external thermocouples are no longer needed. By using a highly accurate on-die temperature sensing circuit, and a fast acting temperature control circuit the processor can rapidly invoke thermal management mechanisms as necessary to control the die temperature. As a result, added thermal design margins can be significantly reduced, and the resulting system performance impact can be minimized if not eliminated.

6.2 Thermal Monitor Implementation

The Intel Xeon processor MP integrates thermal management capability into the processor silicon. This new thermal management capability includes a highly accurate on-die temperature sensing circuit, a signal that indicates the processor has exceeded temperature limits (PROCHOT#), a thermal control circuit (TCC) which can reduce processor temperature by controlling the duty cycle of the processor clocks, and registers to determine the processor thermal status. The processor temperature is determined through an analog thermal sensor circuit comprised of: a diode, a factory calibrated reference current source, and a current comparator (see Figure 6-1). A voltage applied across the diode will induce a current flow that varies with temperature. By comparing this current with the reference current, the processor temperature can be determined. The reference current source corresponds to the diode current when at the maximum permissible processor operating temperature. Each processor is individually calibrated during manufacturing to eliminate any potential manufacturing variations. Once configured, the processor temperature at which the PROCHOT# signal is asserted (trip point) is not re-configurable.

Figure 6-1. Thermal Sense Circuit

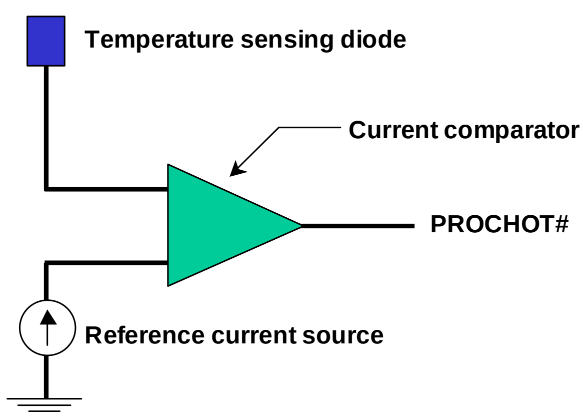

The PROCHOT# signal is used both internal to the processor as well as external to the system. External indication of the processor temperature status is provided through the bus signal PROCHOT#. When the processor temperature is equal to or above the trip point, PROCHOT# is asserted. When the processor temperature is below the trip point, PROCHOT# is de-asserted. Assertion of the PROCHOT# signal is independent of any register settings within the processor and PROCHOT# will be asserted any time the processor die temperature is equal to or exceeds the trip point. The thermal sensor trip point is set to match the maximum specified die temperature.

The Thermal Monitor's thermal control circuit, when active, lowers the processor temperature by reducing the duty cycle of the internal processor clocks. The thermal control circuit portion of the Thermal Monitor must be enabled by the system BIOS for the processor to be operating within specifications. When activated for the Intel Xeon processor MP, the thermal control circuit will turn the processor clocks off and then back on with a predetermined duty cycle.

The actual duty cycle will vary from one product to another. Refer to Figure 6-2 for an illustration. Cycle times are processor speed dependent and will decrease linearly as processor core frequencies increase.

Performance counter registers, status bits in model specific registers (MSR), and the PROCHOT# output pin are available to monitor and control the Thermal Monitor behavior. Details regarding the use of these registers are described in the *Intel® Netburst™ Micro-Architecture BIOS Writers Guide*.

In addition to the Thermal Monitor, the processor clocks can also be modulated via an ACPI register that is implemented as an MSR on the processor core. This is referred to as “on demand mode” clock modulation. See Section 6.3 for additional details.

Figure 6-2. Concept for Clocks Under Thermal Monitor Control

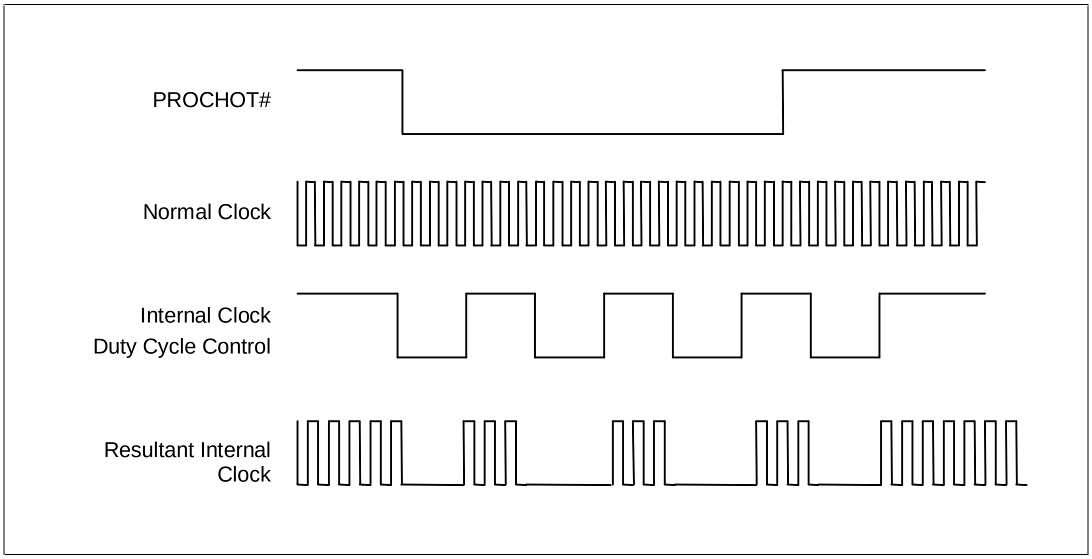

6.3 Operation and Configuration

The Thermal Monitor feature is always enabled. To maintain compatibility with previous generations of processors, which have no integrated thermal management logic, the thermal control circuit portion of Thermal Monitor feature is disabled by default. To utilize the thermal control circuit, either the BIOS enables the feature during the boot process (recommended), or a software driver may enable the feature after the operating system has booted. Refer to the *Intel® Netburst™ Micro-Architecture BIOS Writers Guide* for specific programming details. The thermal control circuit can be configured and monitored in a number of ways. OEMs are expected to enable the thermal control circuit while using various registers and outputs to monitor the processor thermal status. The thermal control circuit is enabled by the BIOS setting a bit in an MSR (model specific register). Enabling the thermal control circuit allows the processor to maintain a safe operating temperature without the need for special software drivers or interrupt handling routines. When the thermal control circuit has been enabled, processor power consumption will be reduced within a few hundred clock cycles after the thermal sensor detects a high temperature (i.e. within a few hundred clock cycles of PROCHOT# assertion). The thermal control circuit and PROCHOT# go inactive once the temperature has been brought back down below the thermal trip point. There is a small hysteresis (~1°C) included in the active to inactive state transition to prevent multiple PROCHOT# transitions around the trip point.

External hardware can monitor PROCHOT# and generate an interrupt whenever there is a transition from active-to-inactive or inactive-to-active. PROCHOT# can also be configured to generate an internal interrupt which would initiate an OEM supplied interrupt service routine. Regardless of the configuration selected, PROCHOT# will always indicate the thermal status of the processor.

For testing purposes, the thermal control circuit may also be activated by setting bits in the ACPI MSRs. The MSRs may be set based on a particular system event (such as an interrupt generated after a system event), or may be set at any time through the operating system or custom driver control thus forcing the thermal control circuit on. This is referred to as “on-demand” mode. Activating the thermal control circuit may be useful for cooling solution investigations or for performance implication studies. When using the MSRs to activate the Thermal Monitor feature, the duty cycle is configurable in steps of 12.5% from 12.5 to 87.5%.

For clock control, the maximum time the clocks will be disabled is ~3 µs. This time period is frequency dependent, and will be shorter on the higher frequency processors. To achieve different duty cycles, the length of time that the clocks are disabled remains constant, and the time period that the clocks are enabled is adjusted to achieve the desired ratio. For example, if the clock disable period is 3µs, and a duty cycle of 1/4 (25%) is selected, the clock on time would be reduced to approximately 1 µs [on time (~1 µs)÷ total cycle time (3 + ~1) µs = 1/4 duty cycle]. Similarly, for a duty cycle of 7/8 (87.5%), the clock on time would be extended to 21 µs [21 ÷ (21 + 3) = 7/8 duty cycle].

In a high temperature situation, if the thermal control circuit and ACPI MSRs (automatic and on-demand modes) are used simultaneously, the fixed duty cycle determined by automatic mode would take precedence.

6.4 System Considerations

The Thermal Monitor feature may be used in a variety of ways, depending upon the system design requirements and capabilities. Intel requires the Thermal Monitor and the thermal control circuit to be enabled for all Intel Xeon processor MP based designs. At a minimum, the thermal control circuit supplies an added level of protection against processor over-temperature failure. Current TDP targets are significantly higher than previous generation Intel® Pentium® III Xeon™ processors (please refer to the appropriate processor datasheet for specific TDP targets). Depending on the system design, this power increase may have a substantial impact on cooling solutions. Larger fans, higher airflow and new heatsinks may be required to maintain proper processor temperatures.

In order to minimize the potential cost impact, system designers are encouraged to take advantage of the Thermal Monitor feature capability. The Thermal Monitor feature allows system designs to address a more realistic TDP, while maintaining protection against the die temperature exceeding the maximum temperature specifications and still maintaining a level of processor performance that is virtually indistinguishable from systems designed to manage maximum power dissipation levels.

Each application program, which is comprised of thousands of processor instructions, will have its own unique power profile, and the profile will have some variability due to loop decisions, I/O activity and interrupts. In general, compute intensive applications with a high cache hit rate will dissipate more processor power than applications that are I/O intensive or have low cache hit rates.

The processor TDP is based on measurements of processor power consumption while running various high power applications. This data was used to determine those applications that are interesting from a power perspective. These applications were then evaluated in a controlled thermal environment to determine their sensitivity to activation of the thermal control circuit. This data was used to derive the TDP targets published in the processor datasheet

A system designed to meet the TDP targets will greatly reduce the probability of real applications causing the thermal control circuit to activate under normal operating conditions. Systems designed for lower power dissipation could be subject to premature activation of the thermal control circuit depending upon ambient air temperature and application power profile. If a designer significantly de-rates the TDP, there is a risk that the Thermal Monitor feature will not be capable of maintaining a safe operating temperature and the processor could shutdown and signal THERMTRIP#. For details regarding the THERMTRIP# signal, refer to Section 6.5.2 or to the processor datasheet.

Figure 6-3 plots processor performance with the Thermal Monitor feature enabled versus system cooling capability. System designers must evaluate the tradeoffs between cooling costs and risk of processor performance loss to determine the optimum configuration for the end user.

Figure 6-3. Processor Performance Versus System Cooling Capability

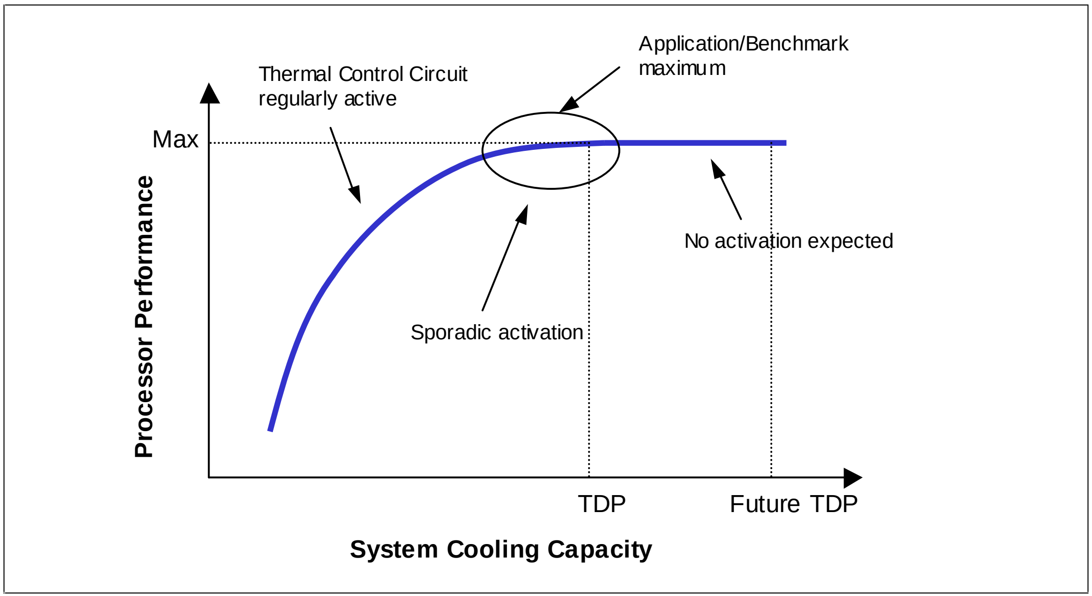

6.4.1 Operating System and Application Software Considerations

The Thermal Monitor feature and its thermal control circuit work seamlessly with any ACPI compliant operating system, provided system BIOS support exists. The Thermal Monitor feature is transparent to application software since the processor bus snooping, ACPI timer and interrupts are active at all times.

6.4.1.1 Operating System Support

Activation of the thermal control circuit during a non-ACPI aware operating system boot process may result in incorrect calibration of software timing loops. The BIOS must disable the thermal control circuit during boot and then the operating system or BIOS must enable the thermal control circuit after the operating system boot process completes. Refer to the *Intel® Netburst™ Micro-Architecture BIOS Writers Guide* for specific programming details. Intel has worked with the major operating system vendors to ensure support for non-execution based operating system calibration loops and ACPI support for the Thermal Monitor feature. Per Microsoft, Microsoft* Windows* 98ES and Microsoft* Windows* 2000 use non-execution based calibration loops and therefore have no issues with the Thermal Monitor feature. When installing Microsoft* Windows NT* 4.0, the user must ensure the PIC-based HAL is used. It is expected that other OS solutions (Linux*, UNIX*, etc) will provide updates to ensure compatibility.

6.5 Legacy Thermal Management Capabilities

In addition to the Thermal Monitor feature, the Intel Xeon processor MP supports the same thermal management features available with the Pentium III Xeon processors. These features include the Thermal Reference Byte located in the Processor Information ROM (PIR), SMBus access to the on-die thermal sensor, and the THERMTRIP# signal for indicating catastrophic thermal failure.

6.5.1 Thermal Sensor

The Intel Xeon processor MP incorporates the SMBus thermal sensor and Thermal Reference Byte features previously enabled for the Pentium III Xeon processor family. Using the SMBus interface, the processor temperature, as measured by the thermal diode, may be read. By averaging this data over long time periods (hours/days vs. min/sec), it may be possible to derive a trend of the processor temperature. Analysis of this information could be useful in detecting changes in the system environment that may require attention.

The processor Thermal Reference Byte is best used for steady state analysis of the processor cooling solution. Attempting to predict the thermal management logic's behavior based on Thermal Reference Byte comparisons is not possible due to the previously described thermal ramp rates. The Thermal Reference Byte is not individually calibrated on the processor, but is generally characterized.

The processor thermal diode should not be relied upon to turn on fans, warn of processor cooling system failure or predict the onset of the thermal control circuit. As mentioned earlier, the processor's high thermal ramp rates make this unfeasible. An illustration of this is as follows. Many thermal sensors report temperatures at a maximum of 8 times per second. Within the 1/8th (0.125 sec) second time period, the temperature is averaged over 1/16th of a second. In a worst-case scenario where the silicon temperature ramps at 50°C/sec, or approximately 6°C/0.125 sec, the processor will be 4.5°C above the temperature reported by the thermal sensor. (Change in diode temperature averaged over 1/16th seconds = 1.5°C, temperature reported 1/16th second later at 1/8th second when the actual processor temperature would be 6°C higher.) The ramp rate error is shown graphically in Figure 6-4.

Figure 6-4. Thermal Sensor Time Delay

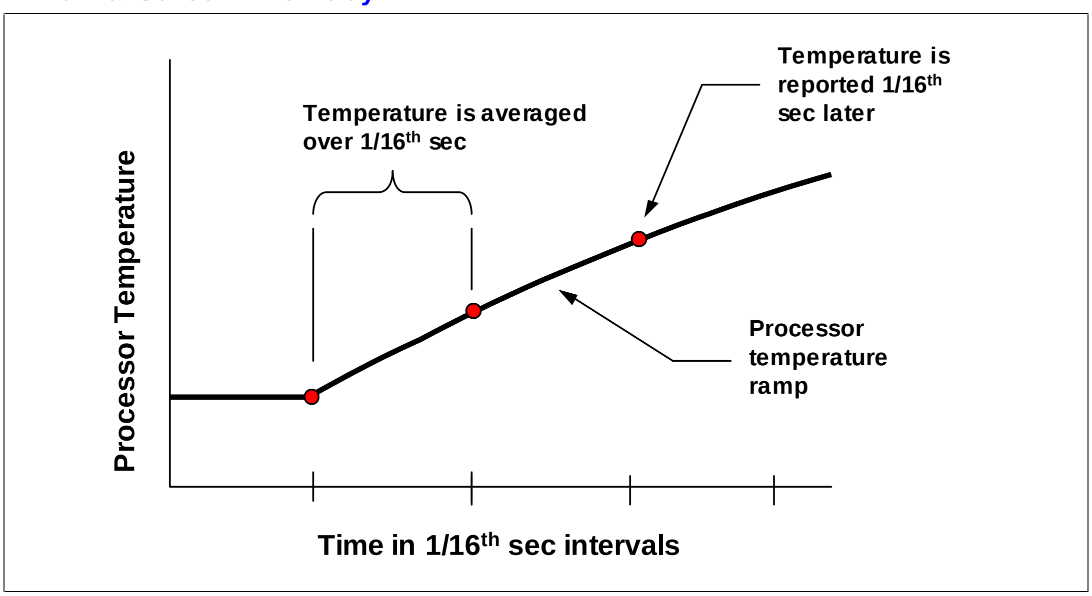

6.5.2 THERMTRIP#

In the event of a catastrophic cooling failure, the processor will automatically shut down when the silicon sensors have reached a numerical value, representing the silicon maximum reliability limits (~135°C). At this point the system bus signal THERMTRIP# will go active and stay active until the processor has cooled down and RESET# has been initiated. THERMTRIP# activation is independent of processor activity and does not generate any bus cycles.

Power must be removed from a processor within 0.5 seconds of THERMTRIP# activation to protect the processor from permanent damage. Because some workstation and server designs that employ multiple processors utilize a shared power plane, power supply sources to all processors must be disabled when any installed processor signals THERMTRIP#. Refer to the processor datasheet or EMTS for timing requirements when designing a circuit to remove power from the processor after THERMTRIP# assertion.

6.5.3 Thermal Measurement Correlation

There are two independent thermal diodes integrated into the Intel Xeon processor MP silicon; one for use with the SMBus thermal sensor and one for the Thermal Monitor feature, which is also used for THERMTRIP#. The Thermal Monitor temperature sensor and the SMBus thermal sensor are independent and isolated devices with no direct correlation to one another. Circuit constraints and their performance requirements prevent the Thermal Monitor sensor and the SMBus thermal sensor from being located at the same place on the silicon. As a result, it is not possible to predict the activation of the thermal control circuit by monitoring the SMBus thermal sensor.

6.6 Cooling System Failure Warning

If desired, the system can be designed to cool the maximum processor power. In this situation, it may be useful to use the PROCHOT# signal as an indication of cooling system failure. Messages could be sent to the system administrator to warn of the cooling failure, while the thermal control circuit would allow the system to continue functioning or allow a graceful system shutdown. If no thermal management action is taken, the silicon temperature may exceed ~135°C causing THERMTRIP# to go active and shut down the processor. Regardless of the system design requirements or cooling solution ability, the Thermal Monitor feature must still be enabled for proper processor operation.

7 Thermal Solution Functional Specifications

This section details the thermal, mechanical, and quality guidelines and requirements for designing Intel Xeon processor MP thermal solutions. Environmental reliability requirements and heatsink critical-to-function dimensions are discussed. Also, the Intel reference thermal solution is presented. With this information, a “third party” could design a thermal solution for the Intel Xeon processor MP.

7.1 Thermal Solution Components

The Intel Xeon processor MP thermal solution shall consist of the following components:

- Heatsink
- Thermal interface material (TIM)
- Heatsink clips
- Retention mechanism (RM)

Figure 7-1 shows the fully assembled and exploded view of the thermal solution components (except for the TIM). The EMI grounding frame shown in Figure 7-1 is an optional component.

Figure 7-1. Exploded View of Thermal Solution Components

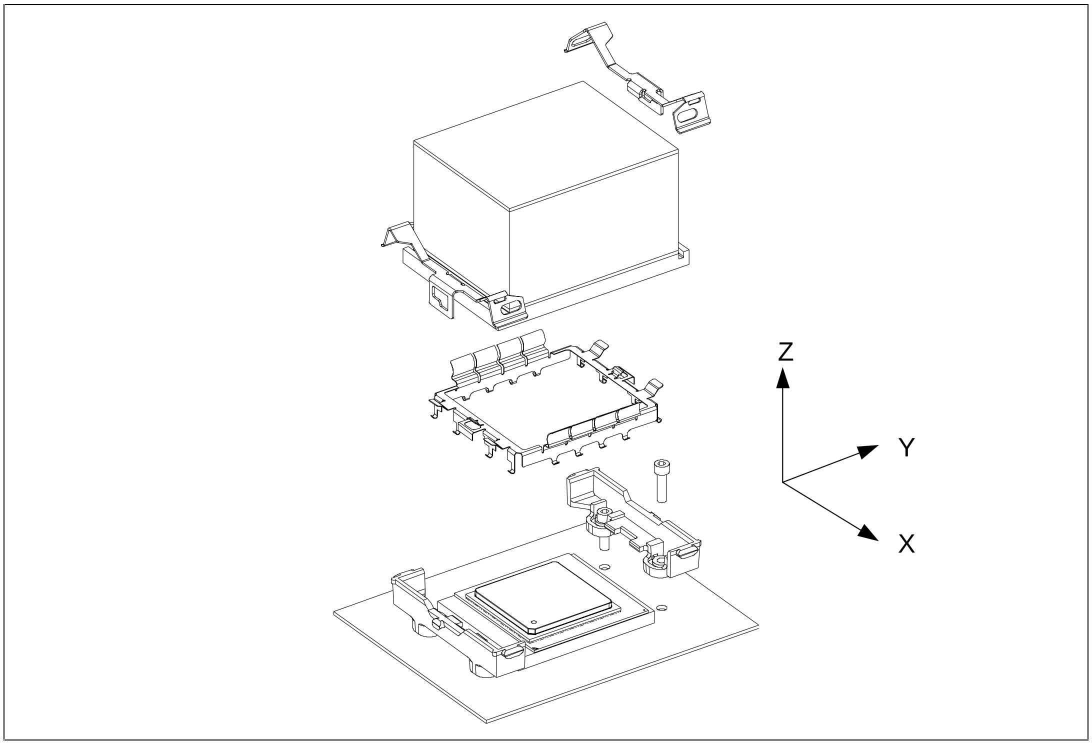

7.2 Design Requirements

7.2.1 Thermal Design Requirements

Thermal solution components should be designed to be in compliance with Intel Xeon processor MP thermal specifications described in the *Intel® Xeon™ Processor MP at 1.40 GHz, 1.50 GHz, and 1.60 GHz Datasheet* and *Intel® Xeon™ Processor MP with up to 2-MB L3 Cache on the 0.13 Micron Process Datasheet* and the design constraints identified in this document. Table 7-1 presents the system design constraints that can be found in typical server and workstation chassis. Workstation chassis may need to implement ducting to achieve this level of airflow.

Table 7-1. System Design Constraints

Maximum Local-Ambient Temperature	Minimum Airflow	Maximum Heatsink Pressure Drop
45°C (113°F)	2.54 m/s (500 lfm)	3.8 mm H_20 (0.15 inches H_20)

For a given chassis, the θ_{CA} requirement is based on the chassis local-ambient characteristics and the processor thermal specifications (T_{case}). The processor thermal solution is required to meet the overall θ_{CA} requirement of the system that it serves. While θ_{CA} is constrained to meet system and processor requirements, θ_{CS} and θ_{SA} are independently constrained. Refer to Section 5.1.1 for definition of thermal resistance relationships, and to Figure 5-1 for a graphical representation of these relationships.

The reference heatsink solutions are evaluated at sea level. However, many system designs must function reliably at high altitude, typically 1,500 m (5,000 ft) or more, and must account for altitude effects on overall thermal performance. Air density changes with altitude may degrade thermal performance for air-cooled thermal solutions. The system designer should account for altitude effects on the overall system thermal design to ensure that temperature specifications for the processor are met at the targeted altitude.

7.2.2 Mechanical Design Requirements

7.2.2.1 Heatsink Critical-to-Function Dimension

Table 7-2 lists the critical-to-function (CTF) dimensions for any heatsink compatible with the enabled retention components. The Intel reference heatsink adheres to these CTF dimensions. Appendix B (Figure B-1 and Figure B-2) provides the drawings detailing the CTF dimensions. Letter references in Table 7-2 highlight the CTF dimensions in Figure B-1 and Figure B-2.

System integrators that choose to utilize alternative heatsink solutions, such as an actively cooled fan heatsink, should refer to the required mechanical envelope provided by the alternative heatsink vendor. The mechanical envelope should include any required clearances above or around the heatsink assembly for adequate airflow into the heatsink fan. Any alternative heatsinks intended for use with the enabled clips and RM should possess the CTF heatsink base dimensions (see Figure B-1) and stay within the heatsink fin volumetric envelope in the X-Y direction (parallel to the heatsink base [see Figure B-2]). Excursions beyond the envelope in the +Z direction may be acceptable depending on specific chassis clearances and alternative heatsink dimensional requirements.

Table 7-2. Critical-to-Function Dimensions

Dimension	Letter	Minimum	Maximum
Location of Clip Attach Groove Far Edge from Heatsink Edge	A	4.57 mm (0.180 in.)	5.08 mm (0.200 in.)
Width of Clip Attach Groove	B	2.03 mm (0.080 in.)	2.54 mm (0.100 in.)
Base Thickness in Zone A	C	6.22 mm (0.245 in.)	6.48 mm (0.255 in.)
Base Length	D	88.6 mm (3.488 in.)	89.2 mm (3.512 in.)
Base Width	E	63.2 mm (2.488 in.)	63.8 mm (2.512 in.)
Base Flatness in Zone B	F	---	0.051 mm/mm (0.002 in./in.)
Width of Clip Attach Area (Zone A)	G	5.08 mm (0.200 in.)	---
Height of Thermal Solution	---	---	50.8 mm (2.000 in.)

7.2.2.2 SSI Dimensional Requirements

The heatsink mechanical envelope must fit within the corresponding system chassis. An effort is underway by the System Server Infrastructure (SSI) initiative to standardize chassis mechanical and electrical specifications for the server and workstation industry. It is advantageous for heatsink designs to be compatible with SSI mechanical constraints in order to reduce redesign costs and ensure compatibility with future system designs. Adopted chassis definitions include the high-end, midrange, and entry electronics bay. A thin electronics bay specification adoption is in process. More information can be found at http://www.ssiforum.org.

Table 7-3 summarizes some of the maximum component height requirements for SSI chassis (all dimensions are superceded by the most current SSI specification):

Table 7-3. SSI Chassis Height Requirements

Chassis Definition	Maximum Component Height
High-End Electronics Bay (Quad IA-32 Processors)	71.12 mm (2.8 in.)
Mid-Range Electronics Bay	152.4 mm (6.0 in.)
Entry-Level Electronics Bay	152.4 mm (6.0 in.) [114.6 mm (4.51 in.) for 3U option]

NOTE: These height requirements do not include baseboard thickness or clearance above chassis.

7.2.2.3 Maximum Heatsink Mass

Heatsinks that attach to the processor via the reference RM should not exceed 450 grams.

7.2.2.4 Heatsink Center of Gravity

The center of gravity of the processor thermal solution should be located over the geometric center of the package. For a heatsink at the recommended maximum mass of 450 grams (1 lbm), the height of the center of gravity should remain below 12.7 mm (0.5 in.) above the bottom of the heatsink base. For heatsinks with mass less than 450 grams, moderate excursions above the recommended limit are acceptable.

7.2.2.5 Heatsink Base Requirements

The flatness of the base shall be maintained at 0.002 in./in maximum at the localized area (Zone B) as shown in Figure B-1. The base plate contains no keying features and thus can be rotated 180 degrees. A heatsink supported by the RM must incorporate two clip attach areas with a minimum width of 0.200 in. (Zone A), as shown in Figure B-1. The heatsink attach clip requirements are presented in Section 7.2.2.7.

7.2.2.6 Thermal Interface Material

A TIM must be applied between the package and the heatsink to ensure thermal conduction. Intel's thermal solution reference designs use either Shin-Etsu* G749 or Shin-Etsu* G751thermal grease. Please refer to the Intel's enabled components list for the latest thermal interfaces used. The materials listed in this document refers to thermal greases. The use of thermal grease in conjunction with high performance heatsink technologies (e.g. copper base folded fin) has been demonstrated to meet Intel thermal performance requirements.

Section 7.5 provides the recommended grease dispense weights to ensure full coverage of the processor IHS. For alternate TIMs, full coverage on Zone B of the heatsink base (see Figure B-1) is recommended to ensure that the entire processor IHS is covered. It will be important to compensate for heatsink to package attach alignment when selecting the proper size of a pre-applied or solid pad-type TIM. If a pre-applied TIM is specified, it may have a protective application tape that must be removed prior to heatsink attach to the processor.

The use of thermally conductive grease as the TIM requires special handling and dispense guidelines. The following guidelines apply strictly to ShinEtsu G749 and G751 thermal grease. Alternate thermal greases may require modifications to either the equipment, process, or application targets depending on the material properties. The use of a semi-automatic dispensing system is recommended for high volume assembly to ensure an accurate amount of grease is dispensed on top of the IHS prior to assembly of the heatsink. A typical dispense system consists of an air pressure and timing controller, a hand held output dispenser, and an actuation foot switch. Thermal grease in cartridge form is required for dispense system compatibility. A precision scale with an accuracy of ±5mg is recommended to measure the correct dispense weight and set the corresponding air pressure and duration. The IHS surface should be free of foreign materials prior to grease dispense

Additional recommendations include recalibrating the dispense controller settings after any two hour pause in grease dispense. The grease should be dispensed just prior to heatsink assembly to prevent any degradation in material performance. Once grease dispense is started, all the grease should be used up or disposed of in appropriate waste containers. (Contact your Environmental, Health, and Safety representative to determine disposal requirements.) Finally, the thermal grease should be verified to be within its recommended shelf life before use.

7.2.2.7 Heatsink Clip Requirements

Heatsink attach clips apply force to the heatsink base to maintain desired pressure on the TIM between the package and the heatsink, and holds the heatsink in place under dynamic loading. The heatsink clip must be designed in a way that minimizes contact with the motherboard surface during clip attach to the RM tab features; the clip should not scrape and/or scratch the motherboard. All surfaces of the clip should be free of sharp edges to prevent injury to any system component or to the person performing the installation.

The Intel reference design heatsink clip will attach to the heatsink base via the grooves at each end of the base, as shown in Figure B-3 and Figure B-4. The reference design heatsink clip is latched to the reference design RM clip tabs, one at each end of the RM.

The clips may be susceptible to deformation during any rework or upgrade procedure where the heatsink assembly is disassembled. Intel's clip design was validated with unused clips that were not subjected to an assembly-disassembly cycle. The system integrator should exercise caution in re-using clips that have experienced one or more assembly-disassembly cycles.

7.2.2.8 Retention Mechanism Requirements

The heatsink retention clips for the Intel Xeon processor MP require a RM. There are no features on the 603-pin socket to directly attach a heatsink. Instead, a RM is used to provide an attachment location for heatsink retention clips.

Intel has determined through extensive mechanical characterization that the use of direct chassis attach of the processor RM can mitigate the risk of mechanical damage to the motherboard, processor, and other surface mounted components in mechanical shock or mechanical drop testing. However, direct chassis attach may not mitigate that risk for all chassis and/or motherboard configurations. Mechanical shock or mechanical drop testing followed by functional and visual quality checks are required for each chassis-motherboard configuration.

Intel's thermal solution reference design uses direct chassis attach of the processor RM. Intel recommends the use of 6-32 [x 3/8-1/2 in.] pan head or round head screw [four each] for direct RM to chassis attach. The screw head must be less than 0.284 in. diameter and less than 0.190 in. height.

7.2.2.9 EMI Ground Frame Requirements

Test results indicate that an EMI grounding frame is not necessary to reduce the electro-magnetic emissions from the Intel Xeon processor MP. As a result, Intel has not enabled tooling for the EMI grounding frame. The grounding frame is an optional component of the Intel reference design and is presented in Appendix B of this document (see Figure B-5).

7.3 Environmental Reliability Requirements

The thermal solution assembly (including all of its components) shall be designed to meet the environmental reliability requirements as outlined in Table 7-4.

Table 7-4. Environmental Reliability Test Conditions

Test	Level
Mechanical Shock	50g, 11 ms, Trapezoidal, 3 drops in each of 5 directions (±X, ±Y, –Z). 30g, 11 ms, Trapezoidal, 3 drops in 1 directions (+Z). See Figure 7-1 for clarification of directions.
Vibration	5-500 Hz, 3.13g RMS, 10 min/axis.
Temperature Cycling	–25°C to 100°C, 10-30°C/min ramp, 15 min dwell, 192 cycles.
Temperature Humidity	95°C, 85% RH, 14 days.
Bake Test	95°C, 16 days, nominal (< 25%) RH.

7.4 Other Requirements

7.4.1 Recycling Recommendation

It is recommended that any plastic component exceeding 25 grams must be recyclable as per the *European Blue Angel* recycling standards.

7.4.2 Safety Requirements

The Intel Xeon processor MP heatsink shall be consistent with the manufacture of units that meet the safety standards:

- UL recognition-approved for flammability at the system level – all mechanical-enabling components must be a minimum UL94V-0 approved.

7.4.3 Agency Requirements

All edges should not be sharp when tested per UL 1439.

7.5 Intel Reference Designs for Enabled Components

The Intel reference heatsink designs includes a copper base, aluminum folded fin or aluminum base, copper/aluminum crimped fin passive heatsink with Shin-Etsu G749 or G751 thermal grease as the TIM. These reference heatsinks are intended for use with Intel Xeon processor MP at all frequencies. .

Table 7-5. Recommended Thermal Grease Dispense Weights

Processor	Recommended Thermal Grease	Dispense Weight (mg)
Intel® Xeon™ processor MP	Shin-Etsu* G749 or G751	600

The reference design accommodates the vertical height of the processor, as specified in the processor datasheet, and the vertical height of the 603-pin socket, as specified in the *603-Pin Socket Design Guidelines*. The Intel reference design heatsink clip will apply a load on the TIM of approximately 25 lbf.

Table 7-6 summarizes the case to ambient thermal resistance, θ_{CA}, of the Intel reference design for the Intel Xeon processor.

Table 7-6. Thermal Resistance Summary of Intel Reference Heatsinks

Reference Heatsink	Thermal Resistance, θ_{ca} (°C/W)
Intel® Xeon™ processor MP	0.32

The figures in Appendix B present the Intel reference design for the heatsink clip, the EMI grounding frame, and the RM. Supplier contact information and mechanical CAD models of these components are available at http://developer.intel.com.

8 Enabled Ducting Solutions

8.1 Processor Wind Tunnel

The processor wind tunnel (PWT) is a low cost, high performance ducting solution that can be configured to fit extended-ATX/SSI/WTX chassis without requiring modifications. The PWT is an efficient application of existing off-the-shelf fan technology. Configurations with multiple fans are possible to enable redundancy in system cooling solutions.

System board layout is critical in implementing PWTs as part of a system thermal solution. Component placement can force special PWT implementations. For instance, some board layouts may cause one of the PWTs to extend beyond the board outline in a DP system. Other layouts may force one of the PWTs to be pressure cooled and the other vacuum cooled in order to retain system airflow direction.

A basic PWT configuration shown in Figure 8-2 and Figure 8-3 will be supplied with the boxed Intel Xeon processor and supported motherboards. The PWT assembly includes three plastic components (a processor shroud, a fan housing, and an end cap), a chamfered fin heatsink, two heatsink spring clips, a 60 x 25 mm fan, and a unique processor RM. The PWT RM provides additional features not currently present on the standard Intel Xeon processor MP RM. These features are necessary to anchor the PWT assembly to the RM.

The PWT cooling system is extremely versatile in that it can be configured to provide vacuum or pressure cooling. In addition, a variety of PWT end caps are available that provide a wide range of thermal options to the system integrator. All PWT end caps are symmetrical; therefore thermal redundancy can be achieved in environments that provide adequate PWT clearance. An optional PWT end cap, the hose adaptor, is currently available but not shipping with the boxed Intel Xeon processor. This component makes it possible for the PWT to provide $T_{ambient\text{-}external}$ via a flexible expandable ducting hose (see Figure 8-4) to provide substantially cooler ambient air directly to the processor heatsink. This option is particularly useful for heavily loaded systems generating high internal ambient temperatures. Additional PWT configurations, board component volumetric keep-outs, and heatsink volumetric keep-ins of the PWT can be viewed in Appendix C of this document.

The PWT makes it possible for many 60 x 25 mm tube-axial fans to meet the airflow requirements of Intel’s high performance processors. Fan supplier information is provided on http://developer.intel.com.

Figure 8-1. Processor Wind Tunnel

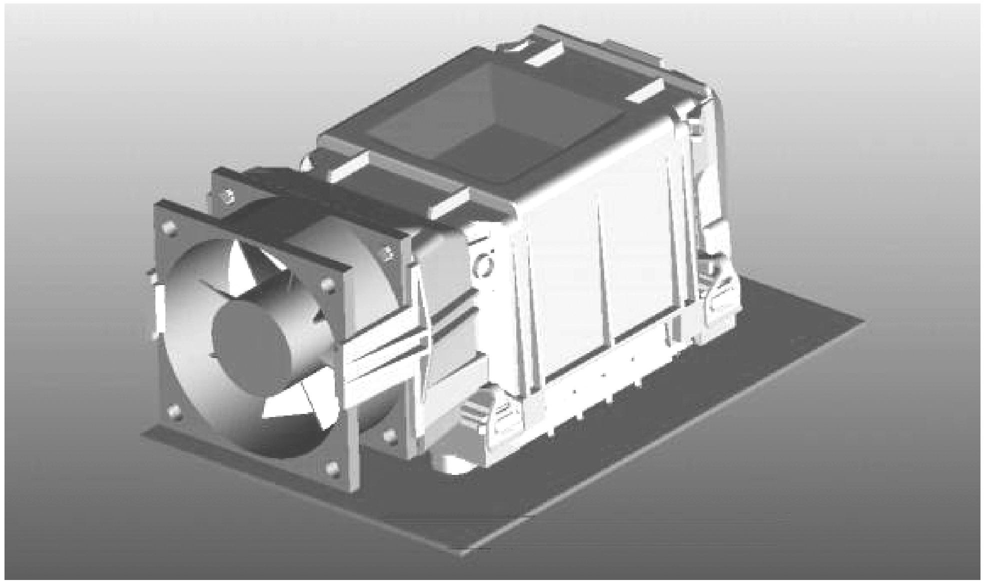

Figure 8-2. PWT Alternate View

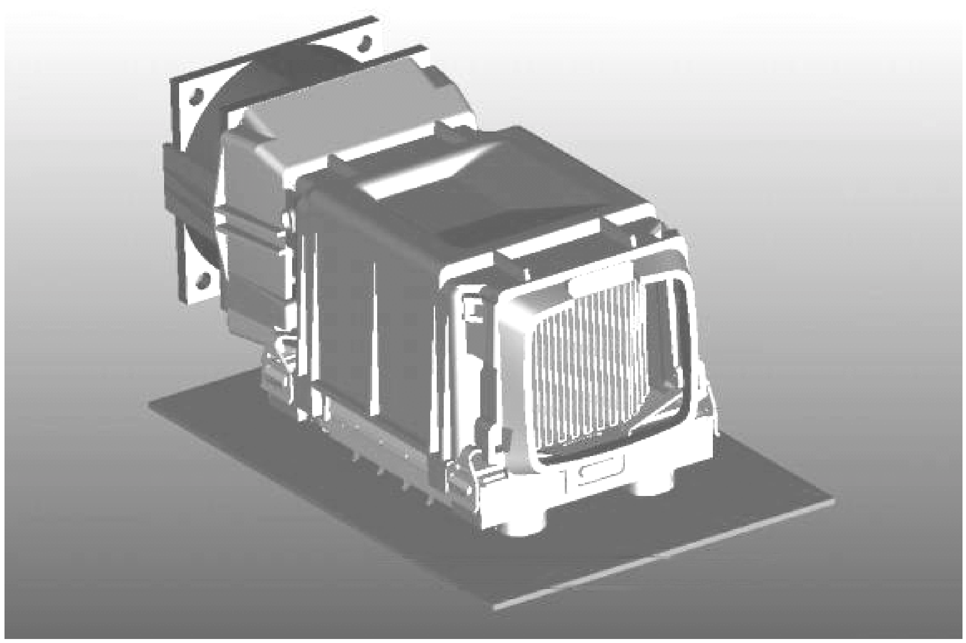

Figure 8-3. PWT with Duct

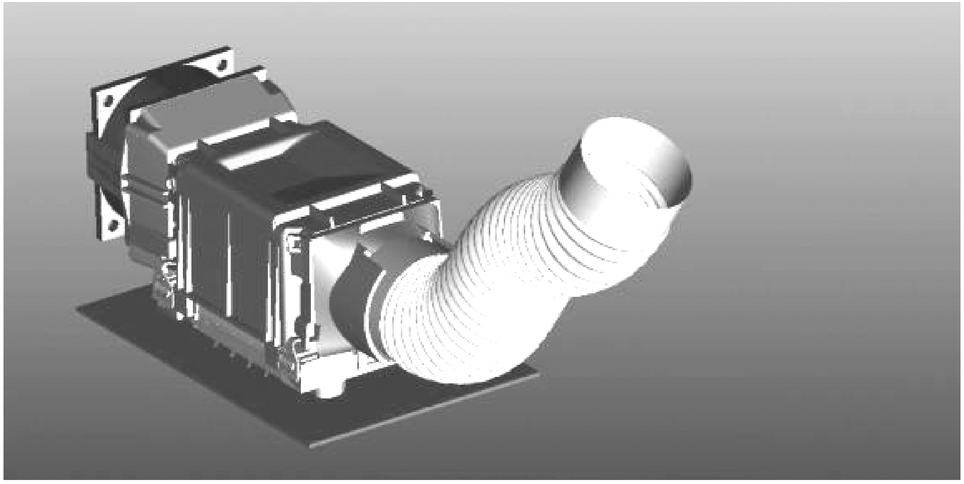

Figure 8-4. PWT with Duct Alternate View

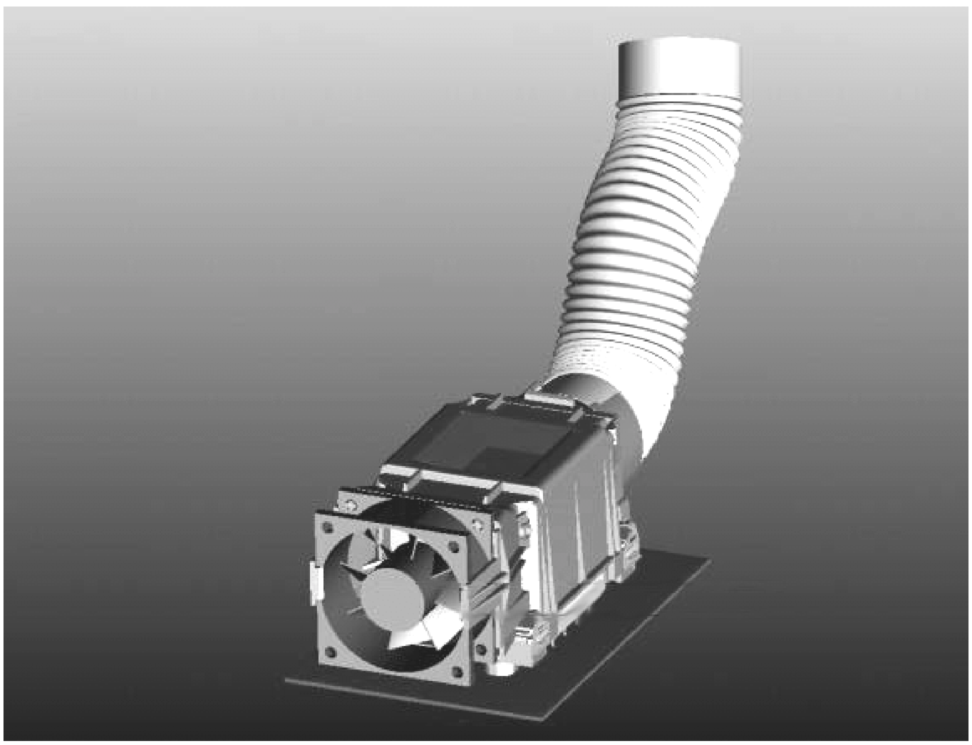

9 Conclusion

As the complexity of today's microprocessors continues to increase, so do the power dissipation requirements. Care must be taken to ensure that the additional power is properly dissipated. Heat can be dissipated using passive heatsinks, fans and/or active cooling devices. Incorporating ducted airflow solutions into the system thermal design can yield additional margin.

The Intel Xeon processor MP has thermal management logic integrated into the processor silicon as well as on-package thermal sensor with SMBus interface. This circuit may be configured to automatically limit the processor temperature through the use of the Thermal Monitor feature. At a factory-calibrated temperature, the processor will periodically stop the internal clocks in order to reduce power consumption and cool down the processor. Various registers and bus signals are available to monitor and control the processor thermal status.

A chassis cooling solution designed to the TDP listed in the processor datasheet will adequately cool the processor to a level where activation of the Thermal Monitor feature is either very rare or non-existent. Various levels of performance versus cooling capacity are available and must be understood before designing a chassis. The OEM has the option to design software to monitor and control the processors thermal capabilities as part of the total system thermal solution.

The size of the heatsink and the output of the fan can be varied to balance size, cost, and space constraints with acoustic noise. This document has presented the conditions and requirements for properly designing a heatsink solution for an Intel Xeon processor MP based system. Properly designed solutions provide adequate cooling to maintain the Intel Xeon processor MP thermal specifications. This is accomplished by providing a low local-ambient temperature and creating a minimal thermal resistance to that local-ambient temperature. Ducting cool external air is highly recommended as a means to lower total cooling solution cost and complexity. By maintaining the processor's case temperature at the values specified in the processor datasheet, a system designer can be confident of proper functionality, performance, and reliability of these processors.

A Designing For Thermal Performance

In designing for thermal performance, the goal is to keep the processor(s) within the operational thermal specifications. Failure to do so will shorten the life of the processor(s) and potentially cause erratic system behavior. The thermal design is required to ensure these operational thermal specifications are maintained. The heat generated by components within the chassis must be removed to provide an adequate operating environment for both the processor and other system components. Moving air through the chassis transports the heat generated by the processor and other system components out of the system, while bringing in air from the external ambient environment.

A.1 Airflow Management

It is important to manage the amount of air that flows within the system as well as how it flows in order to maximize the amount of cool air that flows over the processor. System airflow can be increased by adding one or more fans to the system or by increasing the output (increasing the speed or size) of an existing system fan(s). Managing the local airflow direction using baffles or ducts can also increase local airflow. An important consideration in airflow management is the temperature of the air flowing over the processor(s). Heating effects from chipset, voltage regulators, add-in boards, memory and disk drives greatly reduce the cooling efficiency of this air, as does re-circulation of warm interior air through the system fan. Care must be taken to minimize the heating effects of other system components, and to eliminate warm air circulation.

For example, a clear air path from the system fan(s) to the external system vents will enable the warm air from the processor(s) to be efficiently exhausted out of the system. If no air path exists across the processor(s), the heat generated by the processor(s) will not be removed from the system, resulting in localized heating (“hot spots”) around the processors. Heatsink fin designs should be aligned with the direction of the airflow.

Many multiprocessor system designs will have one processor positioned in front of another processor in the airflow. Without airflow management the second processor will see an increased ambient inlet temperature of about 10-15ºC, depending on the exact layout due to this “thermal shadow” effect. Airflow management, such as ducts and baffles can better distribute cooler air to each of the processors.

A.2 Bypass

Bypass is the distance around the heatsink where air may travel without passing through the fins of the heatsink. A heatsink will have infinite bypass if it is sitting in free space. A duct or other device (such as a hard drive), located 5.1 mm (0.2 in.) from the outer edges of the heatsink, is said to have a bypass of 5.1 mm. A smaller bypass forces more air to pass through the fins of the heatsink, instead of around the heatsink. This is especially important as the heatsink fin density increases. The higher the fin density, the more resistance the heatsink poses to the air and the more likely the air will travel around the heatsink instead of through it unless the bypass is small. Air traveling around the heatsink will have little effect on cooling the processor. Refer to Figure A-1 for an illustration on bypass control.

For multiprocessor configurations, controlling bypass will also require accommodations for partially populated systems. Partially populated systems are systems (in this discussion) which contain less than the maximum number of processors capable of being installed in the system, such as one (1) processor in a four (4) way system. System designers should consider airflow dams or blocking structures to avoid unintentional airflow bypass or pressure drop imbalance.

Figure A-1. Heatsink Bypass Examples

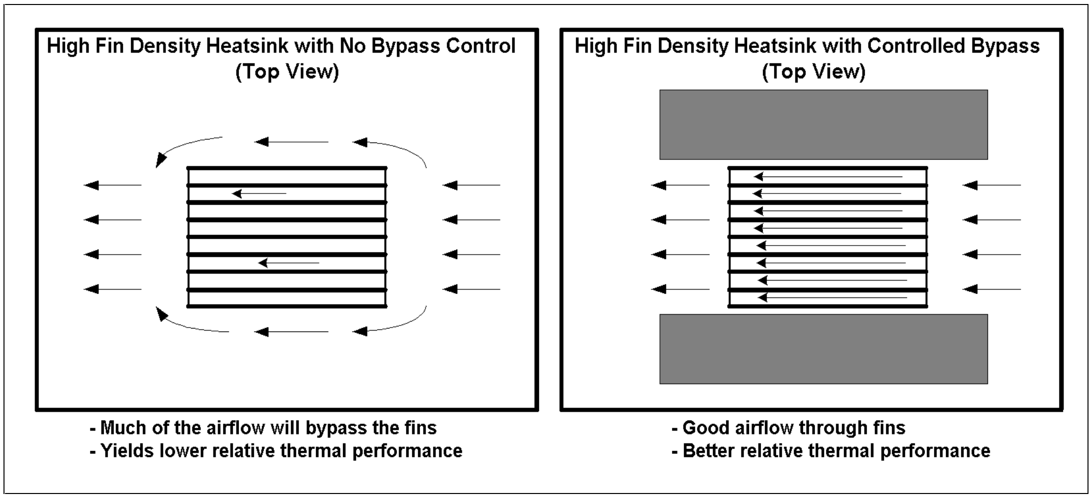

A.3 Heatsink Solutions

One method used to improve thermal performance is to increase the surface area of the device by attaching a metallic heatsink. To maximize the heat transfer, the thermal resistance from the heatsink to the air can be reduced by maximizing the airflow through the heatsink fins as well as by maximizing the surface area of the heatsink itself. As faster processors become available with higher power dissipation, the typical aluminum extruded heatsink may not be sufficient to cool the entire range of TDP. More advanced cooling techniques will likely be required, such as copper base, vapor chamber base, and/or folded or bondedfin heatsink with or without an integrated fan.

Active heatsinks or heatsinks with integrated fans have increased thermal performance efficiencies, due to impinging airflow to the cooling surfaces. The use of active heatsinks in multiprocessor configurations will require special attention to airflow management, intake and exhaust. In addition to cooling solutions, failure mode recovery, such as backup fans, are not well supported in this configuration. Given the complexity in cooling solution and redundancy, many multiprocessor configurations have remained with laminar flow cooling techniques, such as with passive heatsinks.

A.4 Thermal Interface Management

To optimize the heatsink design, it is important to understand the impact of factors related to the interface between the processor and the heatsink base. Specifically, the bond line thickness, interface material area and interface material thermal conductivity should be managed to realize the most effective thermal solution.

A.4.1 Bond Line Management

Any gap between the processor's heat spreader and the heatsink base will impact thermal solution performance. The larger the gap between the two surfaces, the greater the thermal resistance. The thickness of the gap is determined by the flatness of both the heatsink base and the IHS, plus the thickness of the TIM (i.e. thermal grease) used between these two surfaces and the clamping force applied by the heatsink attach clip(s).

A.4.2 Interface Material Area

The size of the contact area between the processor and the heatsink base will impact the thermal resistance. There is, however, a point of diminishing returns. Unrestrained incremental increases in TIM area do not translate to a measurable improvement in thermal performance.

A.4.3 Interface Material Performance

Two factors impact the performance of the interface material between the processor and the heatsink base:

1. Thermal resistance of the material.
2. Wetting/filling characteristics of the material.

Thermal resistance is a description of the ability of the TIM to transfer heat from one surface to another. The higher the thermal resistance, the less efficient the interface is at transferring heat. The thermal resistance of the interface material has a significant impact on thermal performance. The higher the thermal resistance, the higher the temperature drop across the interface and the more efficient the thermal solution (i.e. heatsink) must be to achieve the desired cooling.

The wetting or filling characteristic of the TIM is its ability, under the load applied by the heatsink attach clips, to spread and fill the gap between the processor and the heatsink. Since air is an extremely poor thermal conductor, the more completely the interface material fills the gaps, the lower the temperature drop across the interface. In this case, TIM area also becomes significant; the larger the desired TIM area, the higher the force required to spread the TIM.

Intel has determined through thermal characterization that it may be challenging to meet the thermal performance targets with the use of phase change TIMs. The use of thermal grease in conjunction with high performance heatsink technologies (e.g. copper base folded fin) has been demonstrated to meet Intel thermal performance requirements. The use of thermal grease is recommended. Intel's thermal solution reference designs use Shin-Etsu G751 thermal grease.

A.5 Fans

Fans are needed to move the air through the chassis. The acoustic noise level of a fan is usually directly related to the airflow rate of the fan. Maximum acceptable noise levels may limit the fan output or the number of fans selected for a system. By lowering the ambient temperature at the inlet of the heatsink (T_{LA}), lower airflow across the heatsink and thus slower fan speed can be attained to improve system acoustics. Utilization of ducting or baffles is one method to deliver lower T_{LA} and improve acoustics. See Section A.7.1 for more details regarding ducting

Fan heatsinks are one type of advanced solution that can be used to cool the processor. Due to the concern for reliability, redundancy, and airflow intake/exhaust the current enabled reference design does not include fan heatsink cooling. All of the mechanical features, however, are present for the OEM that wishes to use a fan heatsink assembly.

A.6 Placement

Proper placement of the fans can help ensure that the processor is being properly cooled. Because of the difficulty in building, measuring and modifying a mechanical assembly, models are typically developed and used to simulate a proposed prototype for thermal effectiveness to determine the optimum location for fans and vents within a chassis. Prototype assemblies can also be built and tested to verify if the system components and processor thermal specifications are met.

An intake air fan ideally is centered vertically and placed along one axis with respect to the processor and heatsink. The fan should also be approximately 50.8 mm (2 in.) from the leading edge of the passive heatsink. Figure A-2 shows the fan placement for a typical 4-way layout.

The system fans should be pulling in air from the exterior of the system, and flow directly onto the heatsink. Reduction in preheating of the heatsink airflow results in a degree for degree reduction in the processor case temperature with all other parameters remaining constant.

Figure A-2. Fan Placement and Layout of a Multiprocessor System – Top View

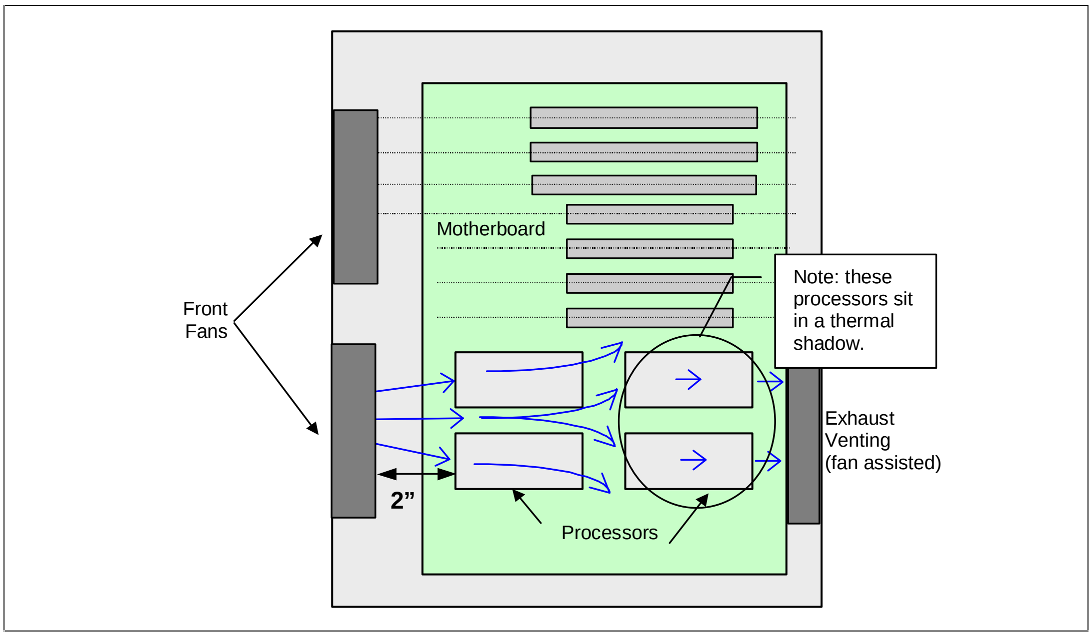

A.6.1 Direction

If the fan(s) are not moving air across the heatsink, then little cooling can occur and the processor may operate above the specified temperature. Two possibilities exist for blowing air across the heatsink of the processor. Air can be blown horizontally, parallel to the baseboard, which blows the air through the length of the heatsink. The air stream can be blown vertically, perpendicular to the baseboard, or down into the heatsink. This may depend on the layout of other components on the board and/or within the chassis. Preferably the intake fan will blow through the heatsink lengthwise because the heatsink fins can be shorter in this case. Both of these factors are considerations when laying out components on the board and in the chassis.

The airflow should be directed with baffles or ducts to flow through the heatsink. This will increase the local flow through the heatsink and may eliminate the need for a second, larger, or higher speed fan.

In dual and quad processor systems the second or successive processor in line could receive pre-heated air. Simulation indicates there could be in excess of a 10ºC rise in air temperature through the first heatsink. This temperature rise makes it necessary to use baffles or ducts to provide cool air to the second processor and exhaust the first processor heated air.

A.6.2 Size And Quantity

It does not necessarily hold true that the larger the fan the more air it blows. A small blower using ducting might direct more air over the heatsink than a large fan blowing non-directed air over the heatsink. The following provide some guidelines for size and quantity of the fan(s).

The fan should be a minimum of 80 mm (3.15 in.) square, with a minimum airflow of approximately 500 linear feet per minute (LFM) at the inlet to the heatsink. Ideally two (2) fans should be used. The intake air fan would blow directly into the processor and heatsink assembly, while a second fan, possibly the power supply would exhaust the air out of the system. For server products, multiple, redundant fans must also be considered for high-reliability systems. These recommendations may not apply if special system solutions, such as fan ducts, are used.

A.6.3 Venting

Intake ports should be placed at the front (user side) of the system to avoid any re-circulation that can occur from the rear of a system with little wall clearance. Vent location placement should include consideration for cooling of processor and peripherals (drives and add-in cards). Intake venting directly in front of the intake fan is the most optimal location. The ideal design will provide airflow directly over the processor heatsink.

A.6.3.1 Placement

Exhaust venting in conjunction with the power supply exhaust fan is usually sufficient for smaller systems. However, depending on the number, location and types of add-in cards, exhaust venting may be necessary near the adapter cards. This should be modeled or prototyped for the optimum thermal potential. Hence, a system should be modeled for the worst-case; i.e. all expansion slots should be occupied with typical add-in options.

A.6.3.2 Area and/or Size

The area and/or size of the intake vents should depend upon the size and shape of the fan(s). Adequate air volume must be obtained and thus will require adequate sized vents. Intake vents should be located in front of the intake fan(s) and adjacent to the drive bays. Venting should be approximately 50% to 60% open in the electro-magnetic interference (EMI) containment area. Outside the EMI containment area, the open percentage can be greater if needed for aesthetic appeal (i.e. bezel/cosmetics). Caution should be exercised that venting is not excessive or poorly placed which can cause re-circulation of warm exhaust air.

A.6.3.3 Vent Shape

Round, staggered pattern openings are best for EMI containment, acoustics and airflow balance.

A.7 Alternative Cooling Solutions

In addition to passive heatsinks, fan heatsinks and system fans, other solutions exist for cooling integrated circuit devices. For example, ducted blowers, heat pipes and liquid cooling are all capable of dissipating additional heat. Due to their varying attributes, each of these solutions may be appropriate for a particular system implementation. More information on this topic can be located on Intel's web site at http://developer.intel.com.

A.7.1 Ducting

Ducts can be designed to isolate the processor(s) from the effects of system heating (such as add-in cards), and to maximize the processor cooling temperature budget. Air provided by a fan or blower can be channeled directly over the processor and heatsink, or split into multiple paths to cool multiple processors. This method can also be employed to provide some level of redundancy in a system requiring redundant capabilities for fault tolerance. This is accomplished by channeling air from two or more fans through the same path across a processor. Each fan, or each set of fans, must be designed to provide sufficient cooling in the event that the other has failed.

A.7.1.1 Ducting Placement

When ducting is to be used, it should direct the airflow evenly from the fan through the length of the heatsink. This should be accomplished, if possible, with smooth, gradual turns, as this will enhance the airflow characteristics. Sharp turns in ducting should be avoided. They increase friction and drag and will greatly reduce the volume of air reaching the processor heatsink. Duct placement should include cooling considerations for other heat dissipating devices. The system designer may want to include auxiliary cooling for other devices, such as voltage regulator modules, memory, etc., in the processor ducting design. If auxiliary cooling is not implemented, the system designer should ensure any duct design does not impede airflow to these devices.

A.8 System Components

A.8.1 Placement

Peripherals such as CD-ROMs, floppy drives, hard drives, VRM's, etc. can be placed to take advantage of a fan's movement of ambient air (by placing them near intake or exhaust fans or venting). Some add-in cards have a low tolerance for temperature rise. These components should be placed near additional venting if they are downstream of the processor to minimize an increase in their ambient temperature.

A.8.2 Power

Some types of drives, such as floppy drives, do not dissipate much heat, while others (e.g. read/write CD-ROM drives, SCSI drives) dissipate a great deal of heat. These hotter components should be placed near fans and/or venting whenever possible. The same can be said for some types of add-in cards. Some PCI cards are very low wattage (approximately 5 W) while others can be as high as 25W, per the PCI specification. AGP graphics devices can dissipate up to 25W per AGP revision 2.0 specifications while AGP Pro50 devices dissipate 25-50W and AGP Pro110 devices dissipate 50-110W per AGP Pro revision 1.1a specifications. Great care should be taken to ensure that these cards have sufficient cooling, while not adversely affecting the processor cooling.

A.9 Voltage Regulation Module Considerations

A.9.1 Airflow and Local-ambient Temperature

Voltage regulation modules (VRMs) must also be considered in system cooling solutions. Because proper power delivery to the processor demands that the VRM be placed very close to the processors, local-ambient temperature for the VRM may be affected by the heating of the nearby processors. A typical VRM requires 400 LFM of airflow to remain within temperature specifications, assuming a local-ambient air temperature of 60ºC. VRM temperature specifications typically specify a maximum temperature at a physical location on the module. In addition to maximum operating temperature, VRM reliability (or MTBF) is very dependent on the operating temperature. System thermal modeling should include the VRM(s) in the simulation to ensure they remain within specifications and reliability expectations. VRM current delivery capability and reliability may be improved with a comprehensive system cooling design.

B Mechanical Drawings

The following table lists the mechanical drawings included in this section.

Figure B-1. Heatsink Base Dimensions

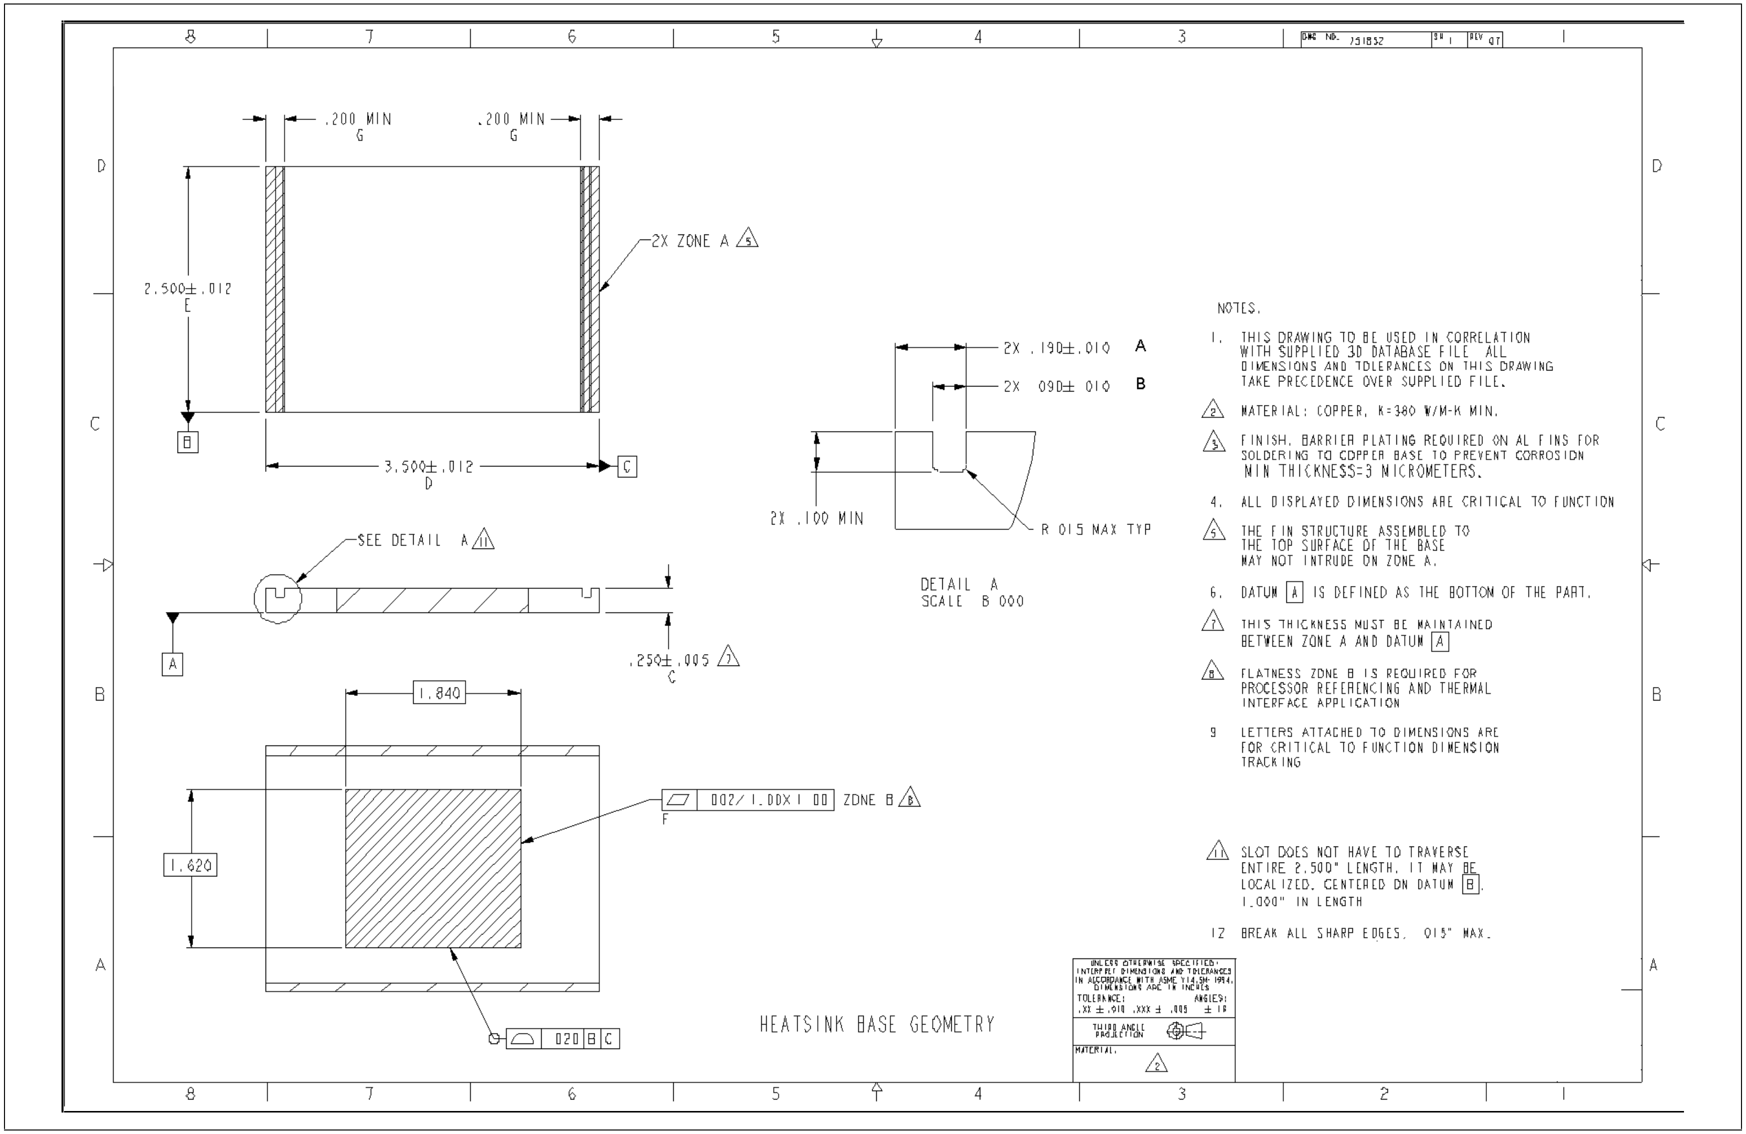

Figure B-2. Heatsink Volumetric Keep-in Zone

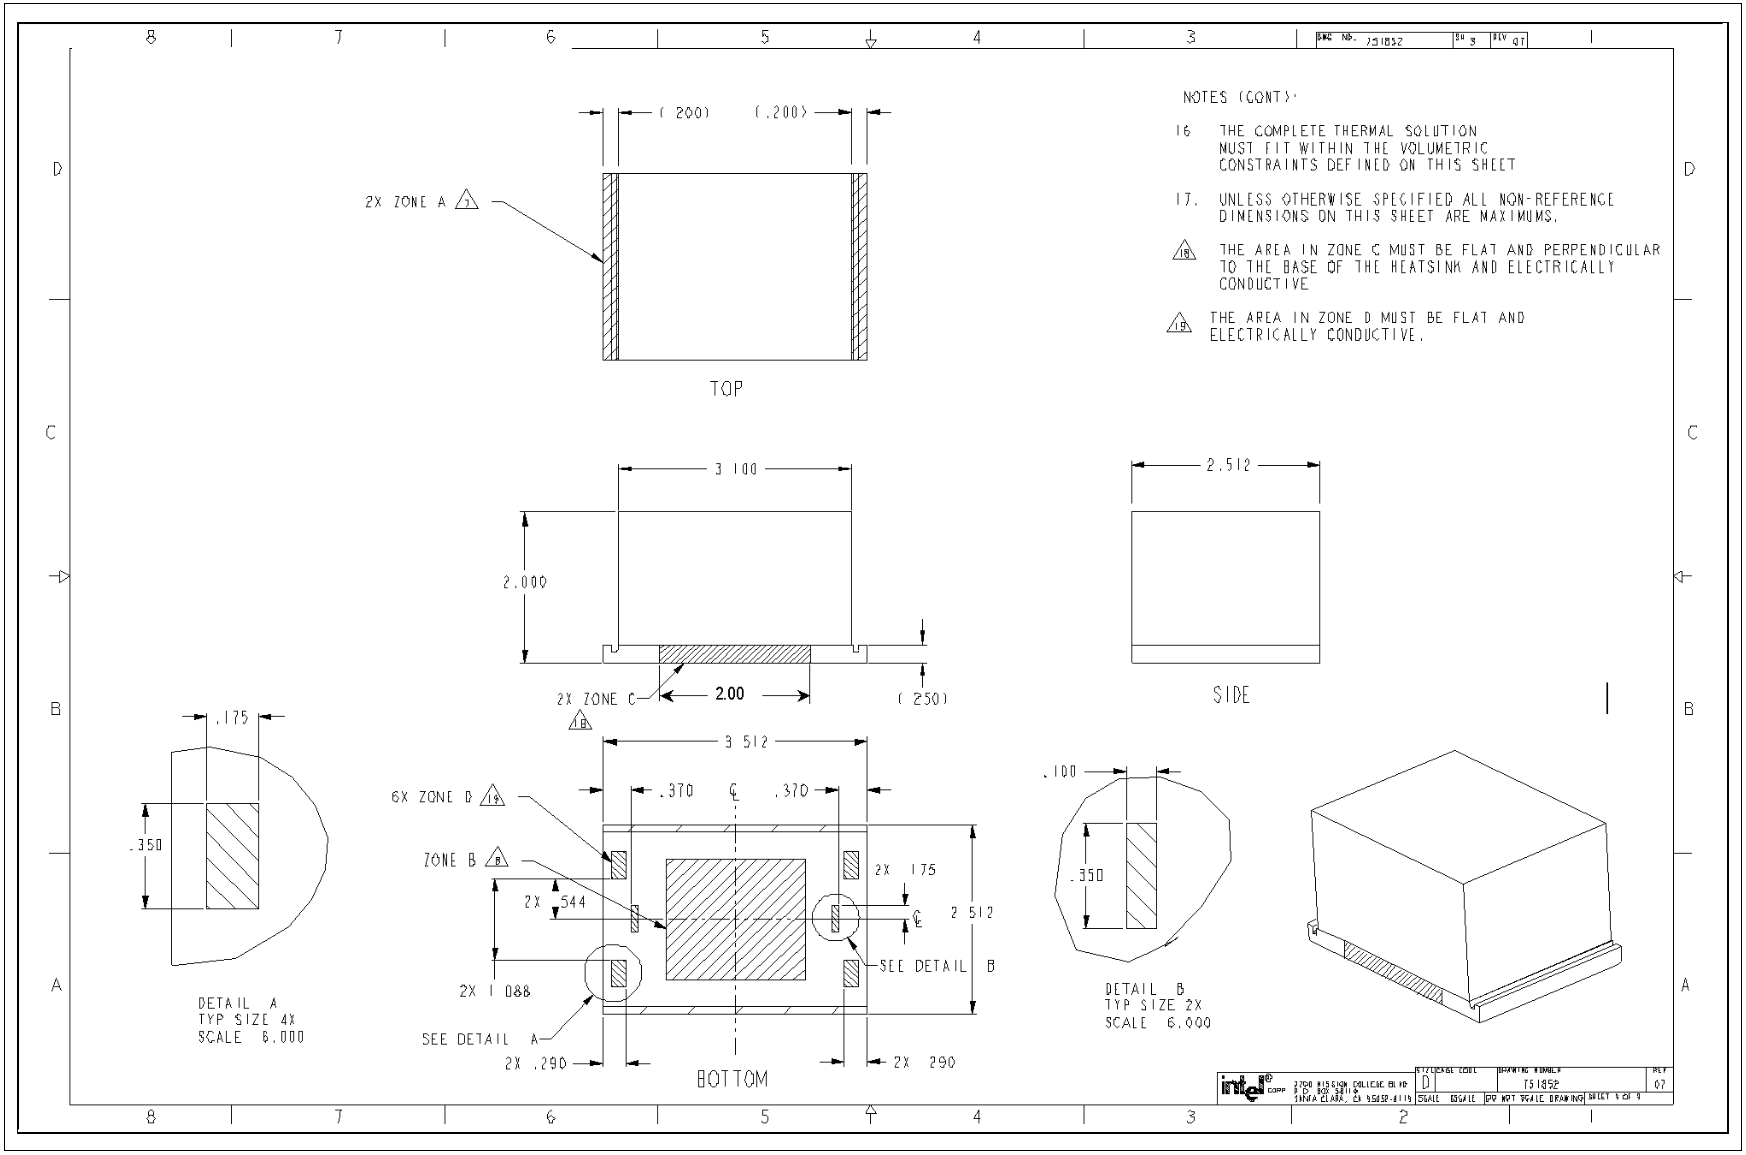

Figure B-3. Enabled Heatsink Clip (Sheet 1 of 2)

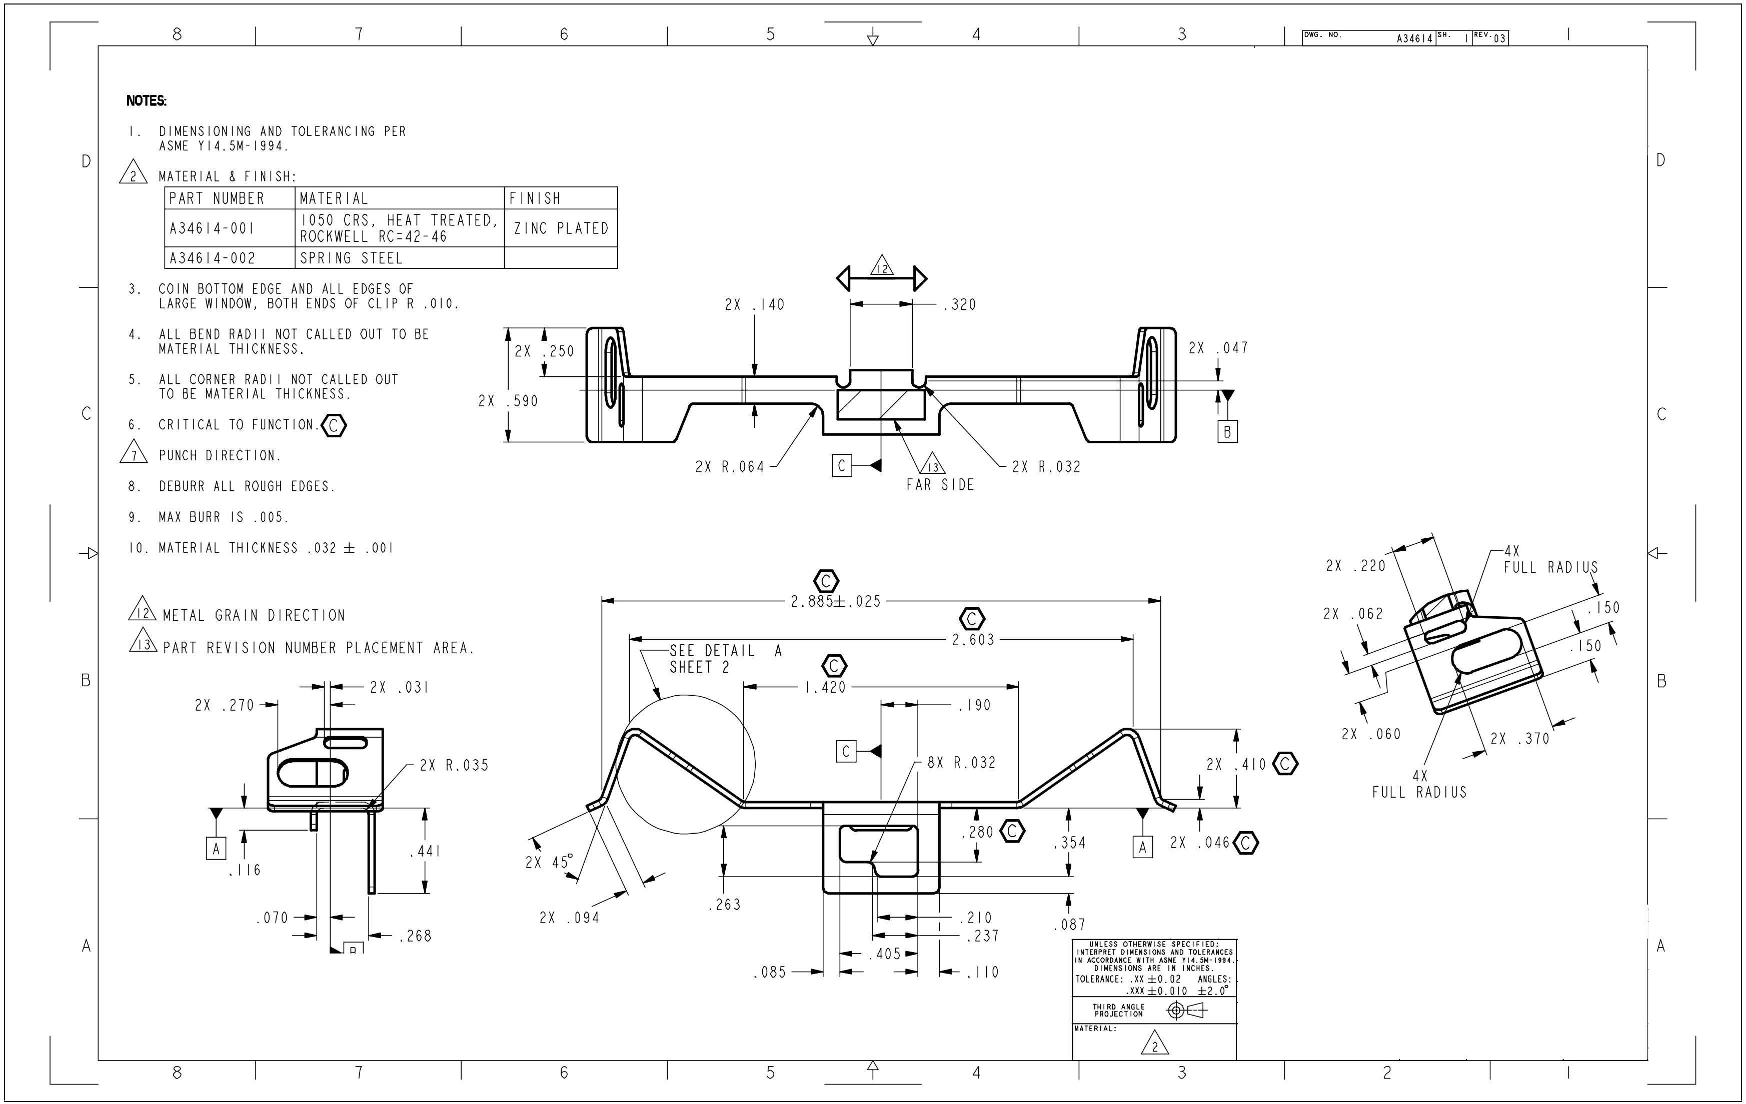

Figure B-4. Enabled Heatsink Clip (Sheet 2 of 2)

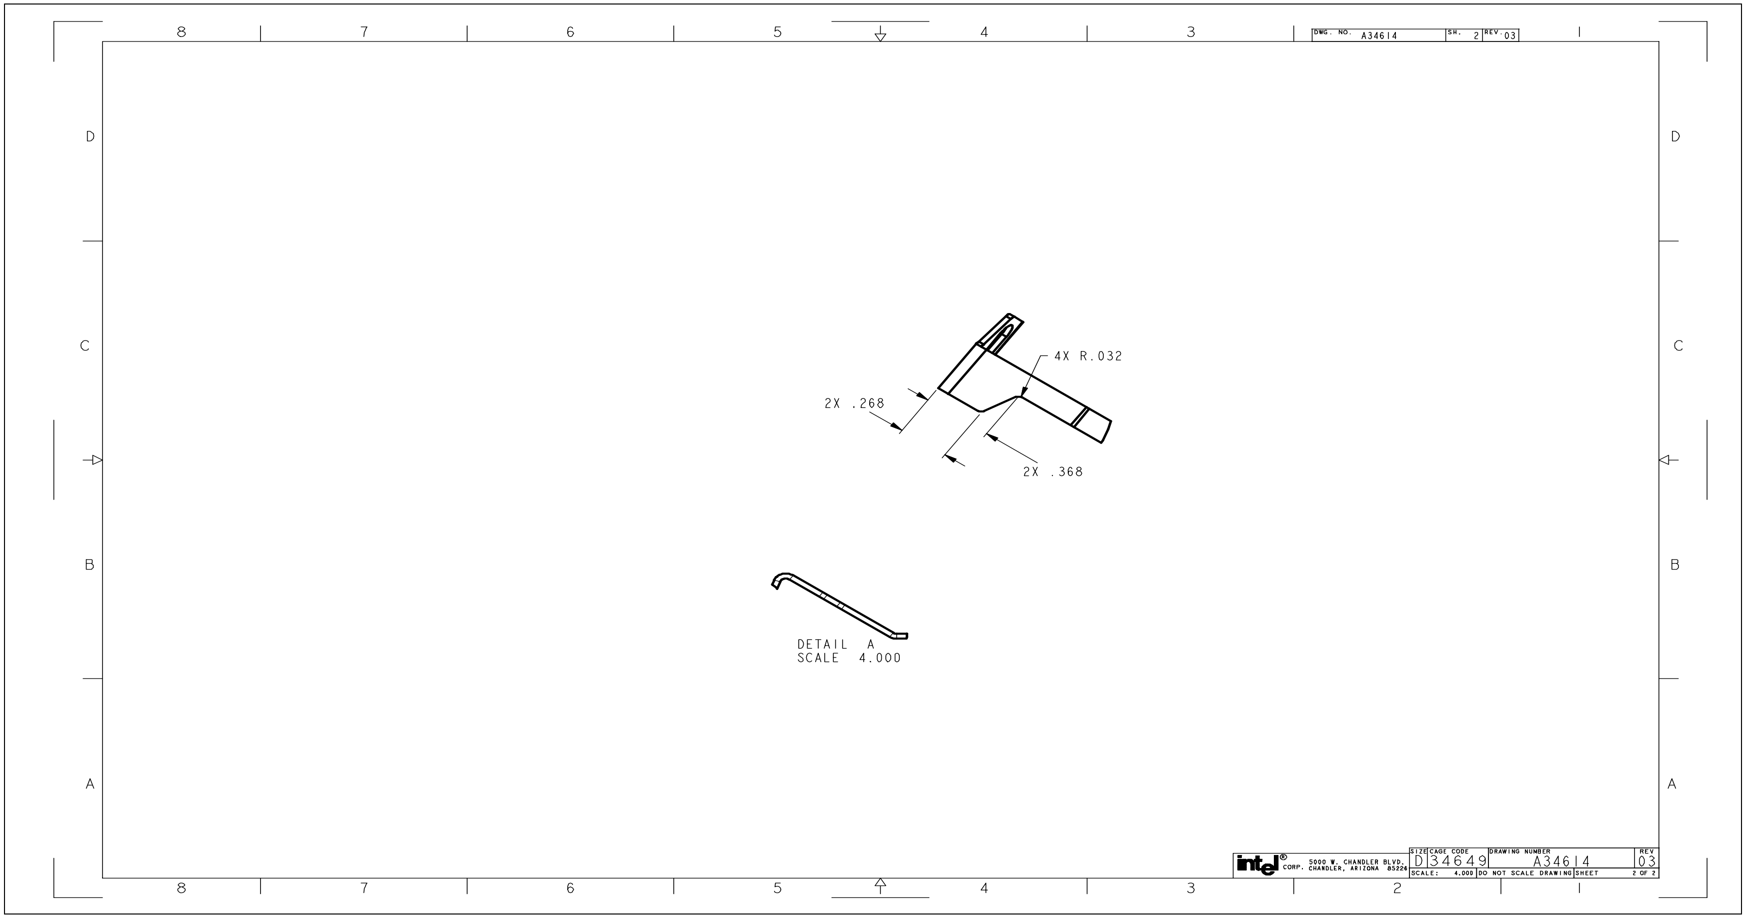

Figure B-5. Enabled EMI Shield

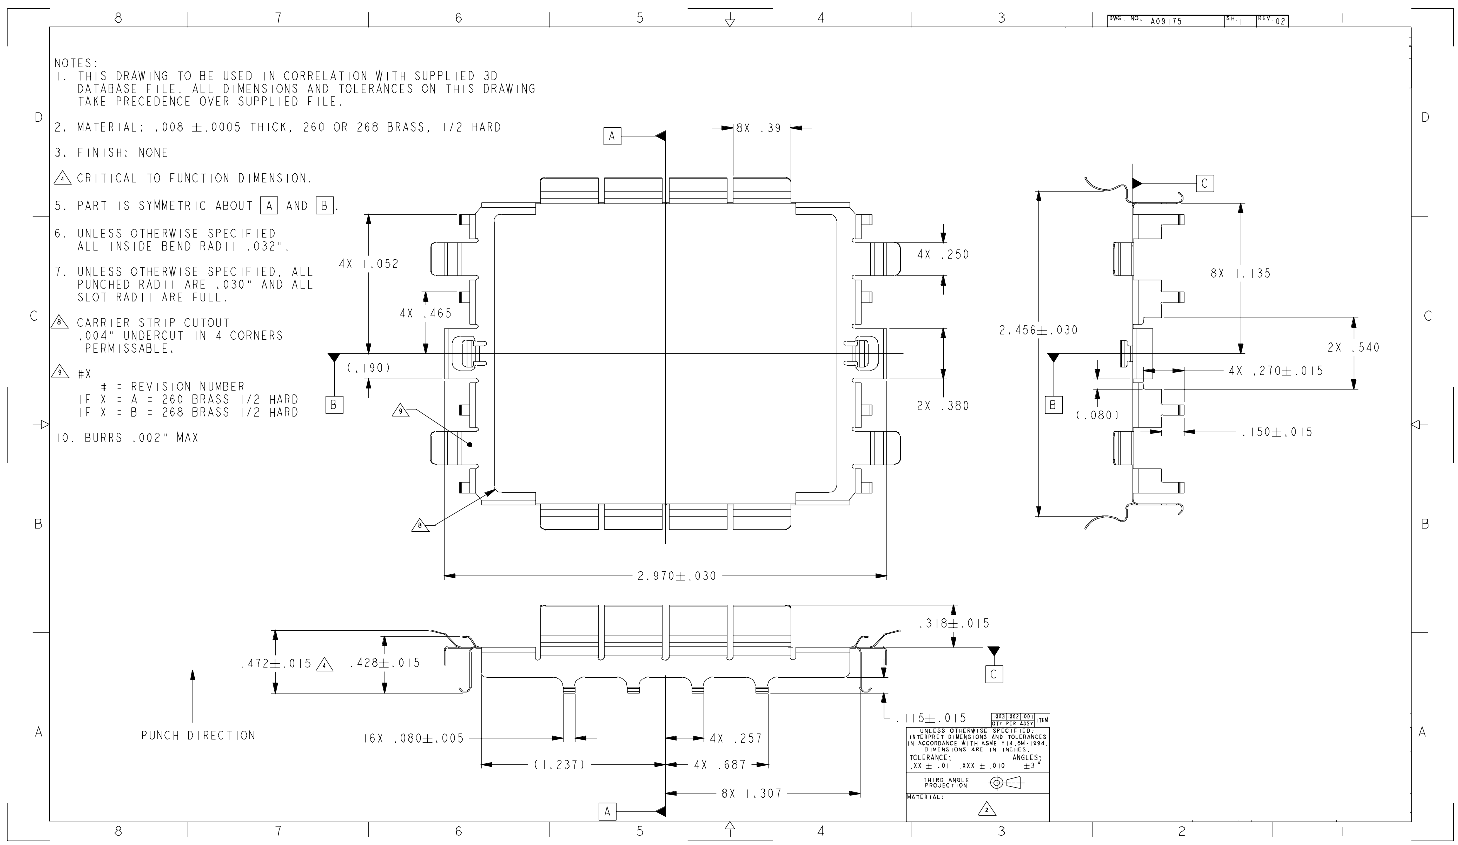

Figure B-6. Enabled Retention Mechanism (Sheet 1 of 4)

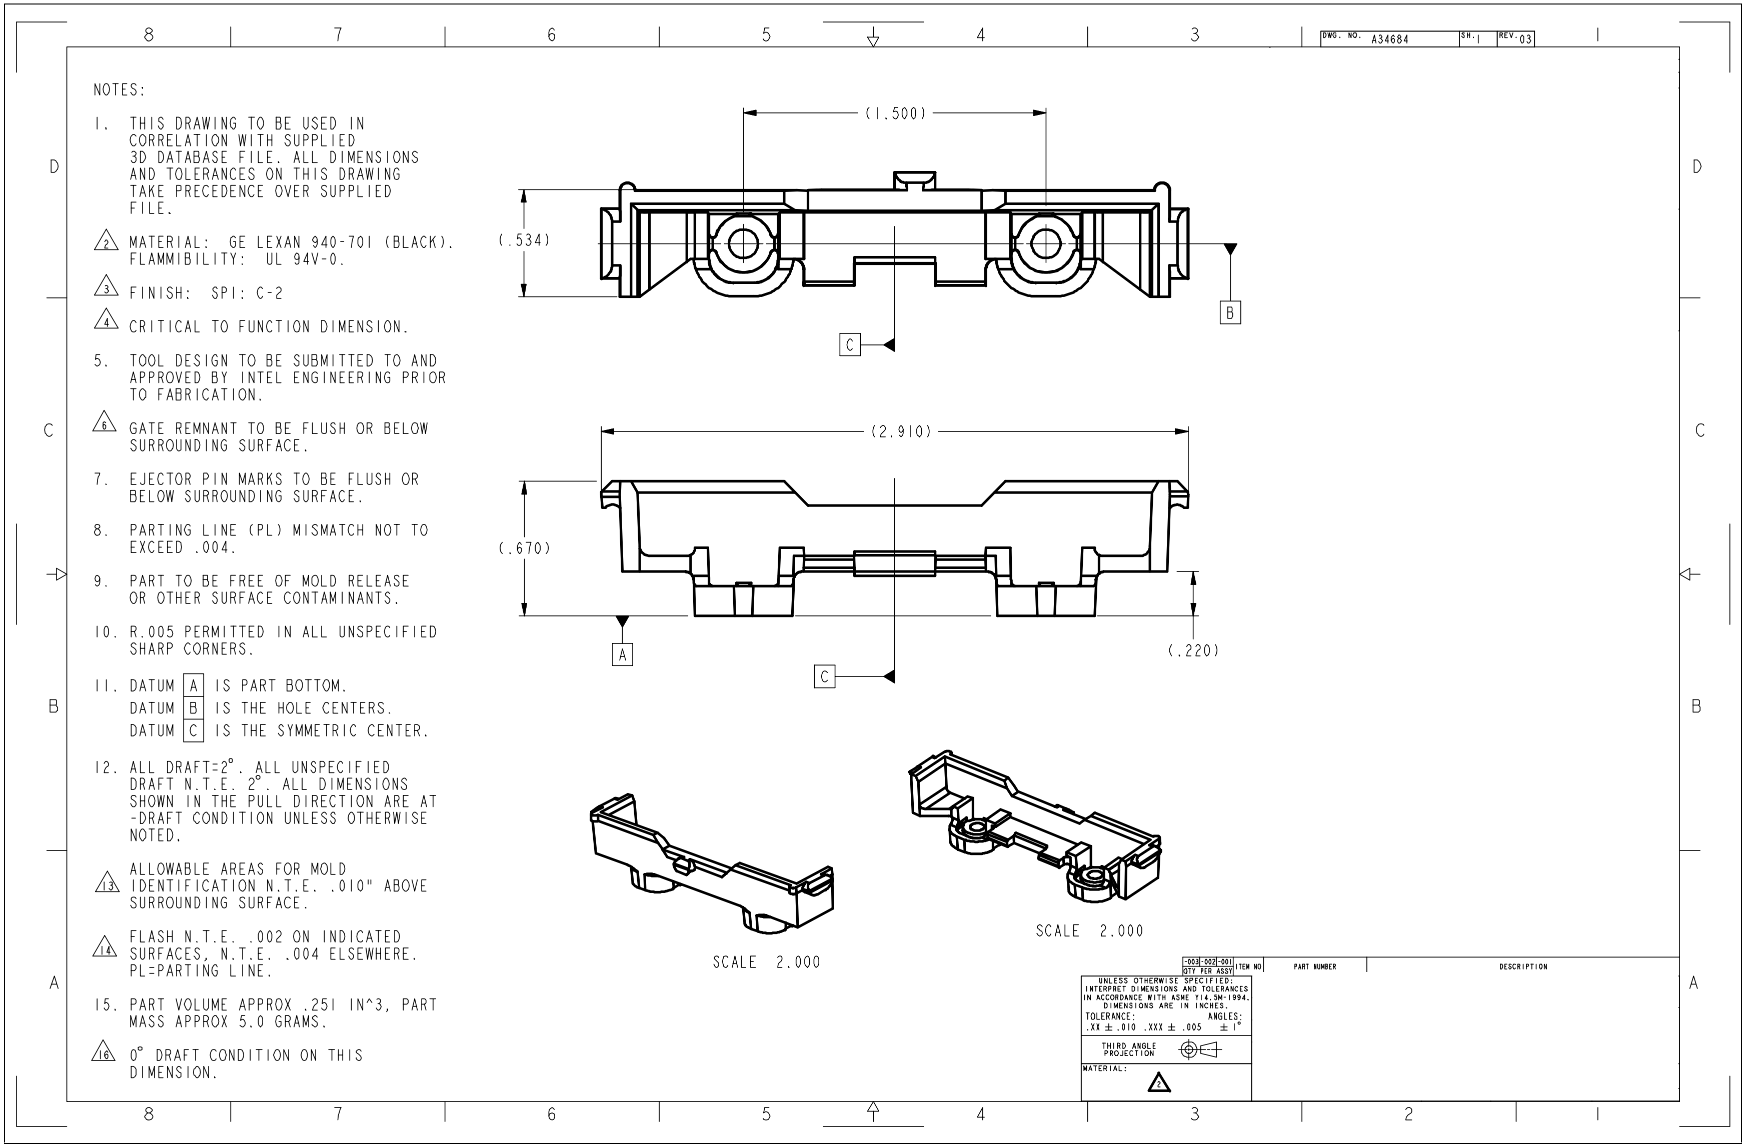

Figure B-7. Enabled Retention Mechanism (Sheet 2 of 4)

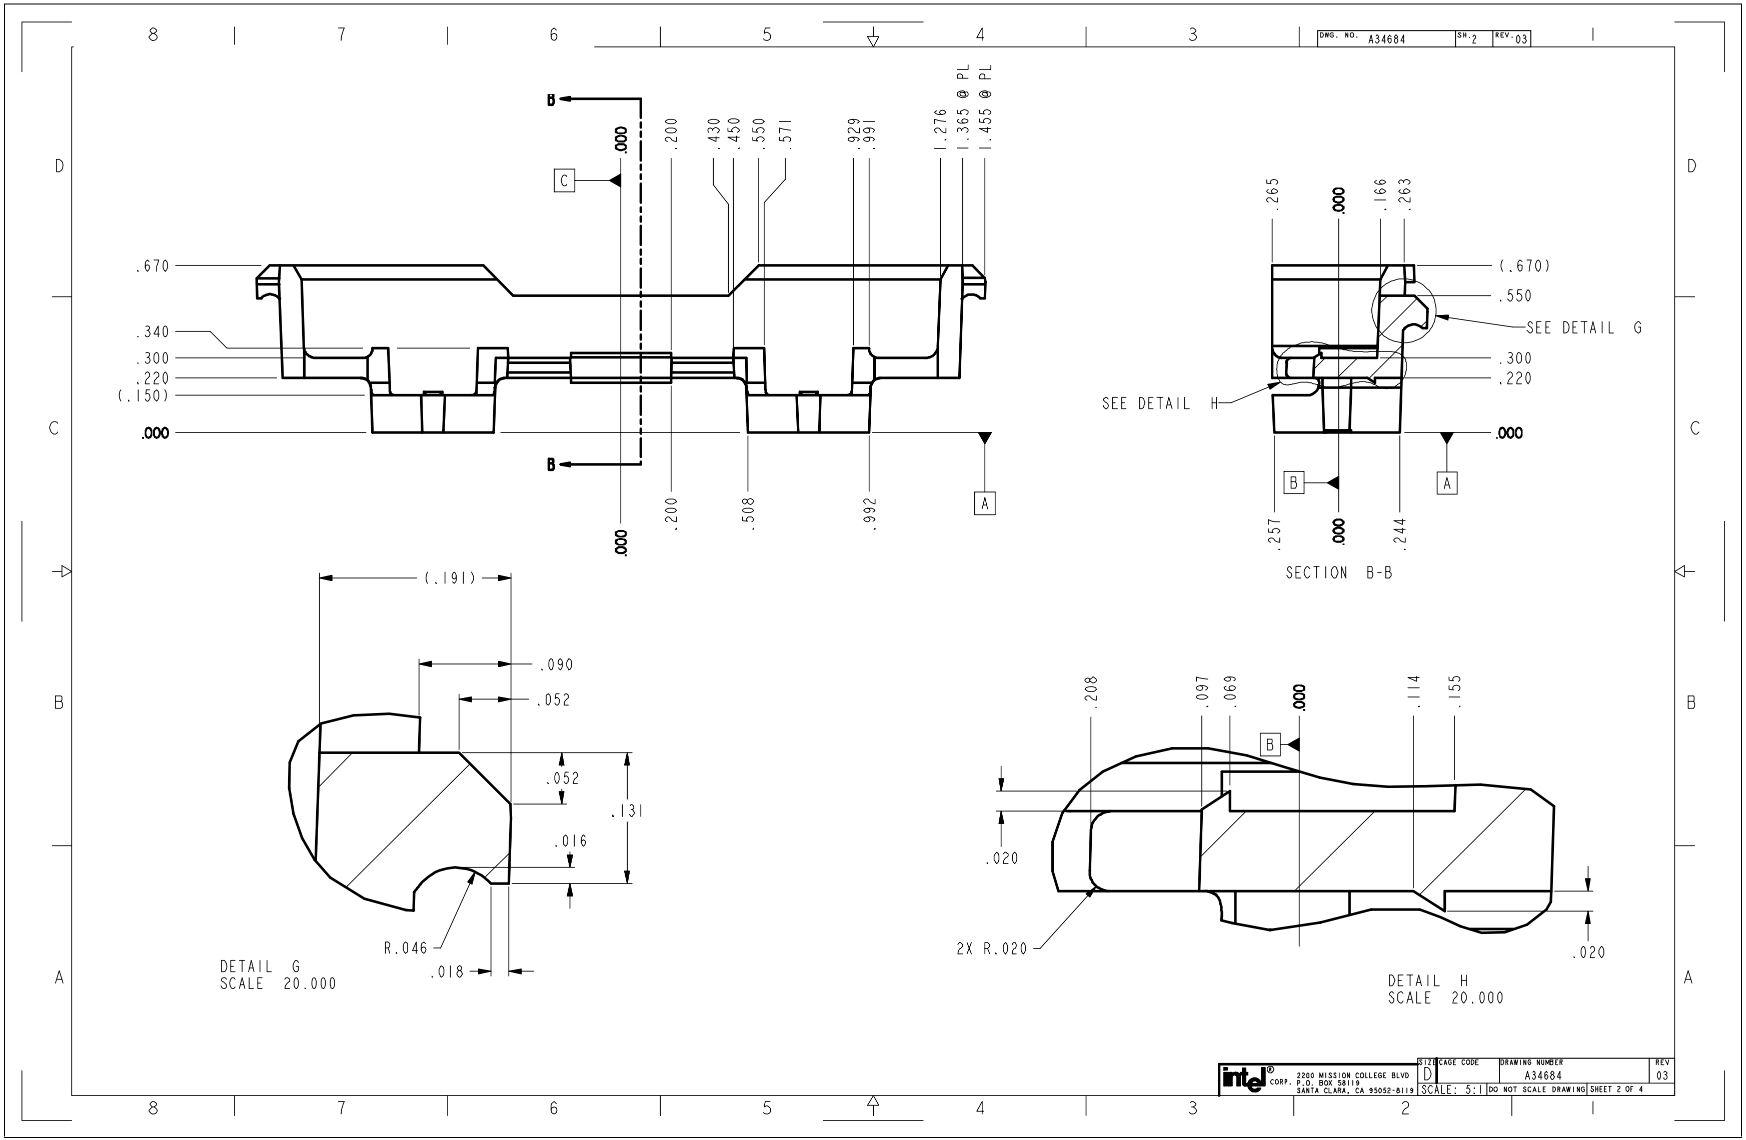

Figure B-8. Enabled Retention Mechanism (Sheet 3 of 4)

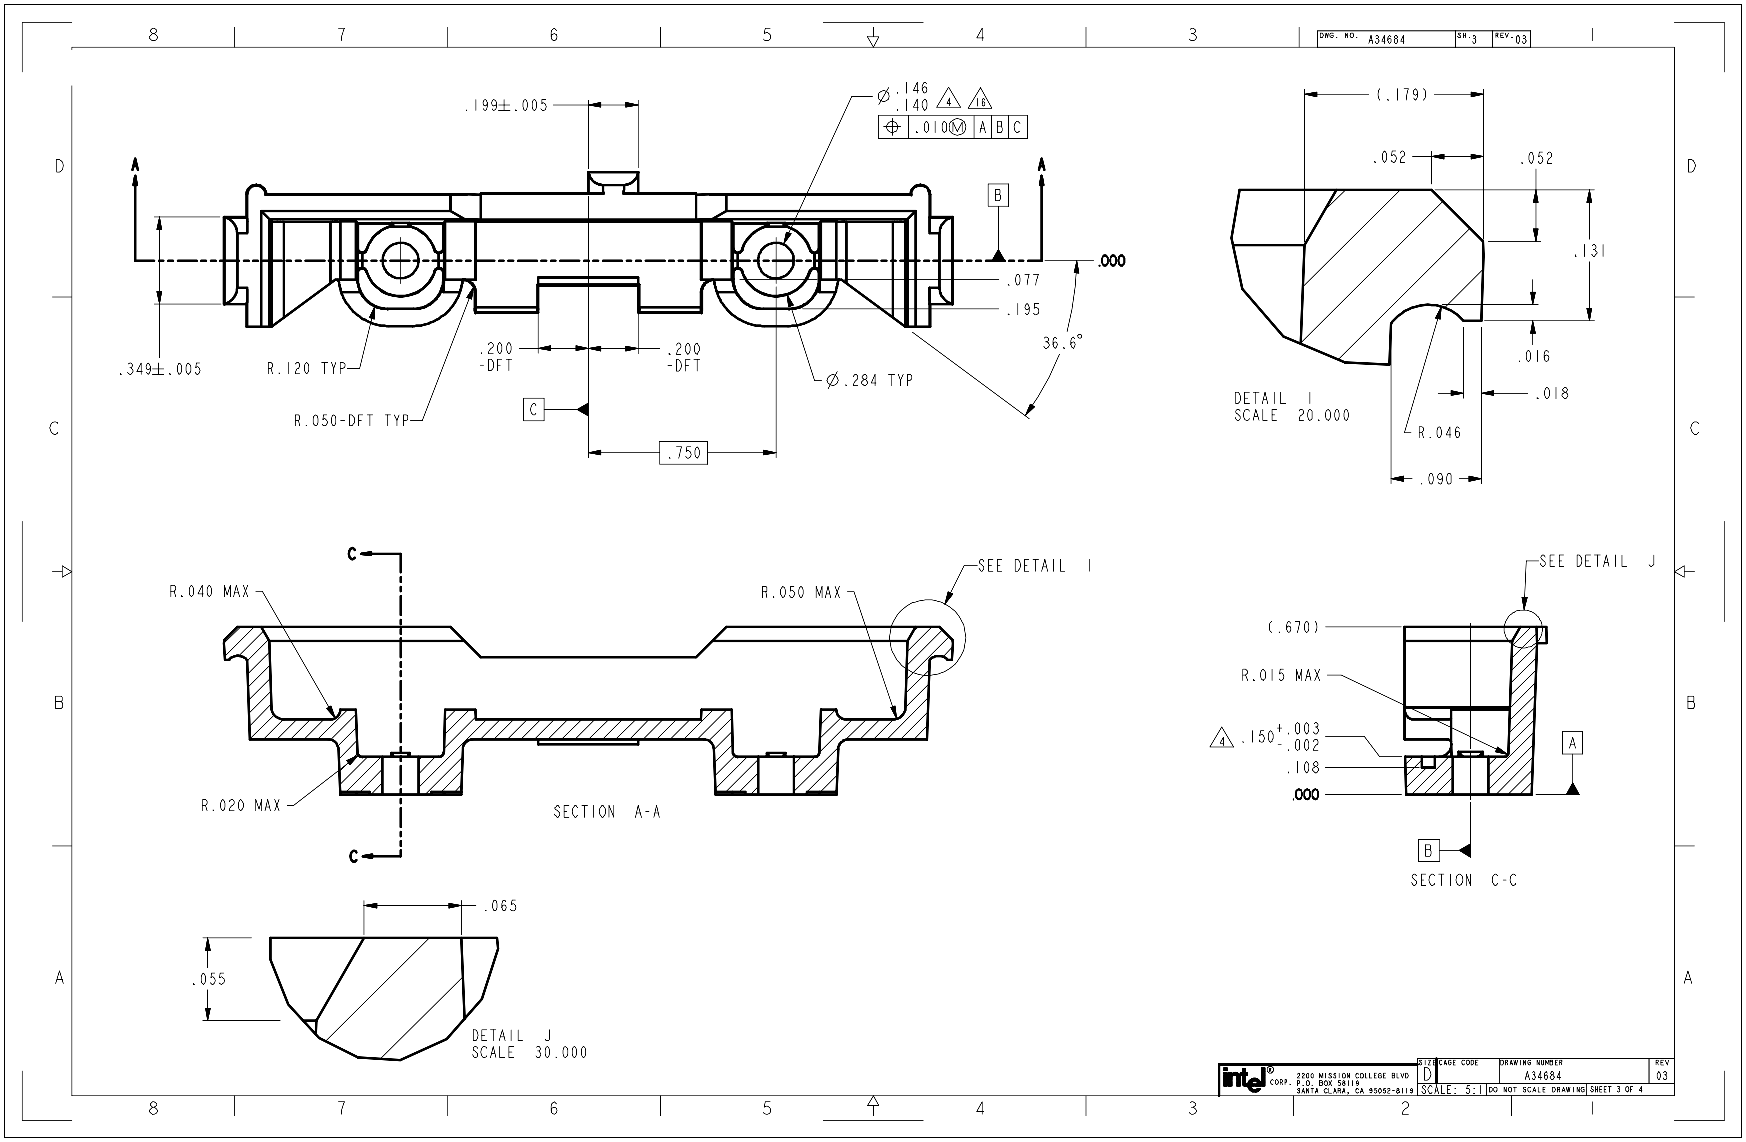

Figure B-9. Enabled Retention Mechanism (Sheet 4 of 4)

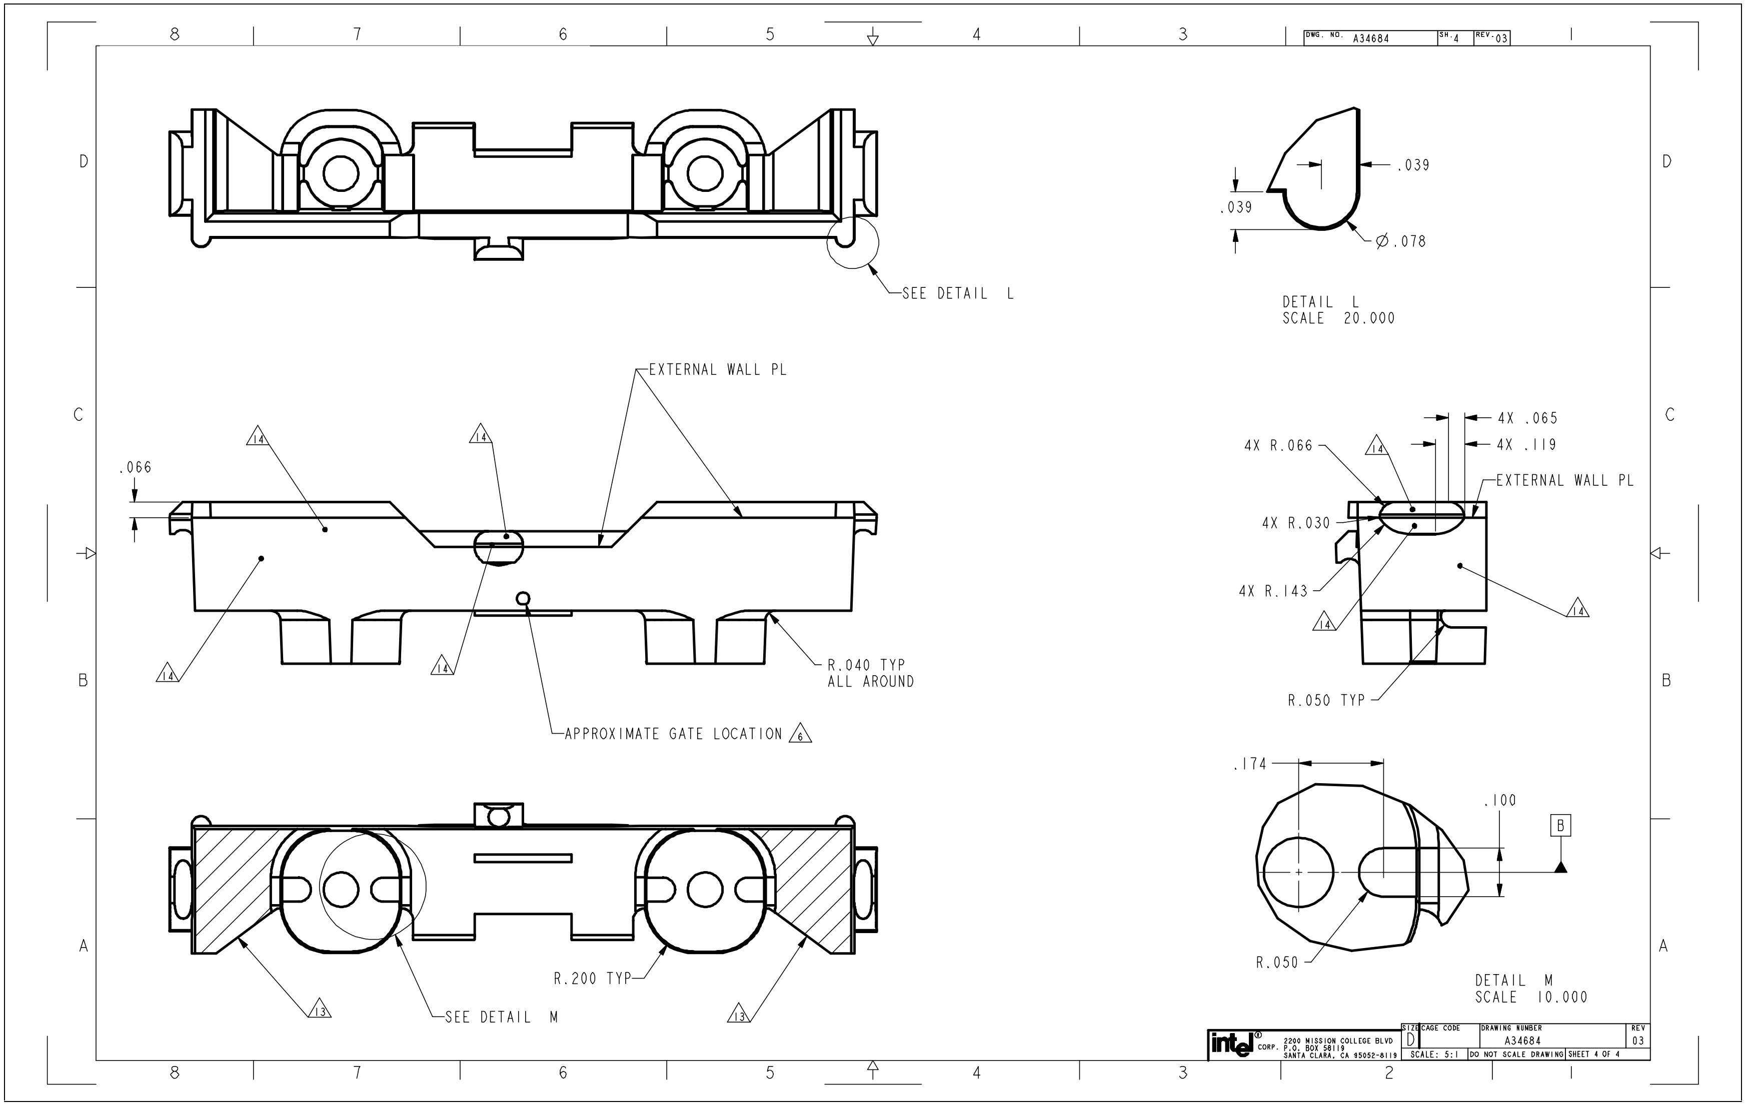

C Processor Wind Tunnel Drawings

The following table lists the mechanical drawings included for the PWT.

Figure C-1. Chamfered Heatsink Keep-In

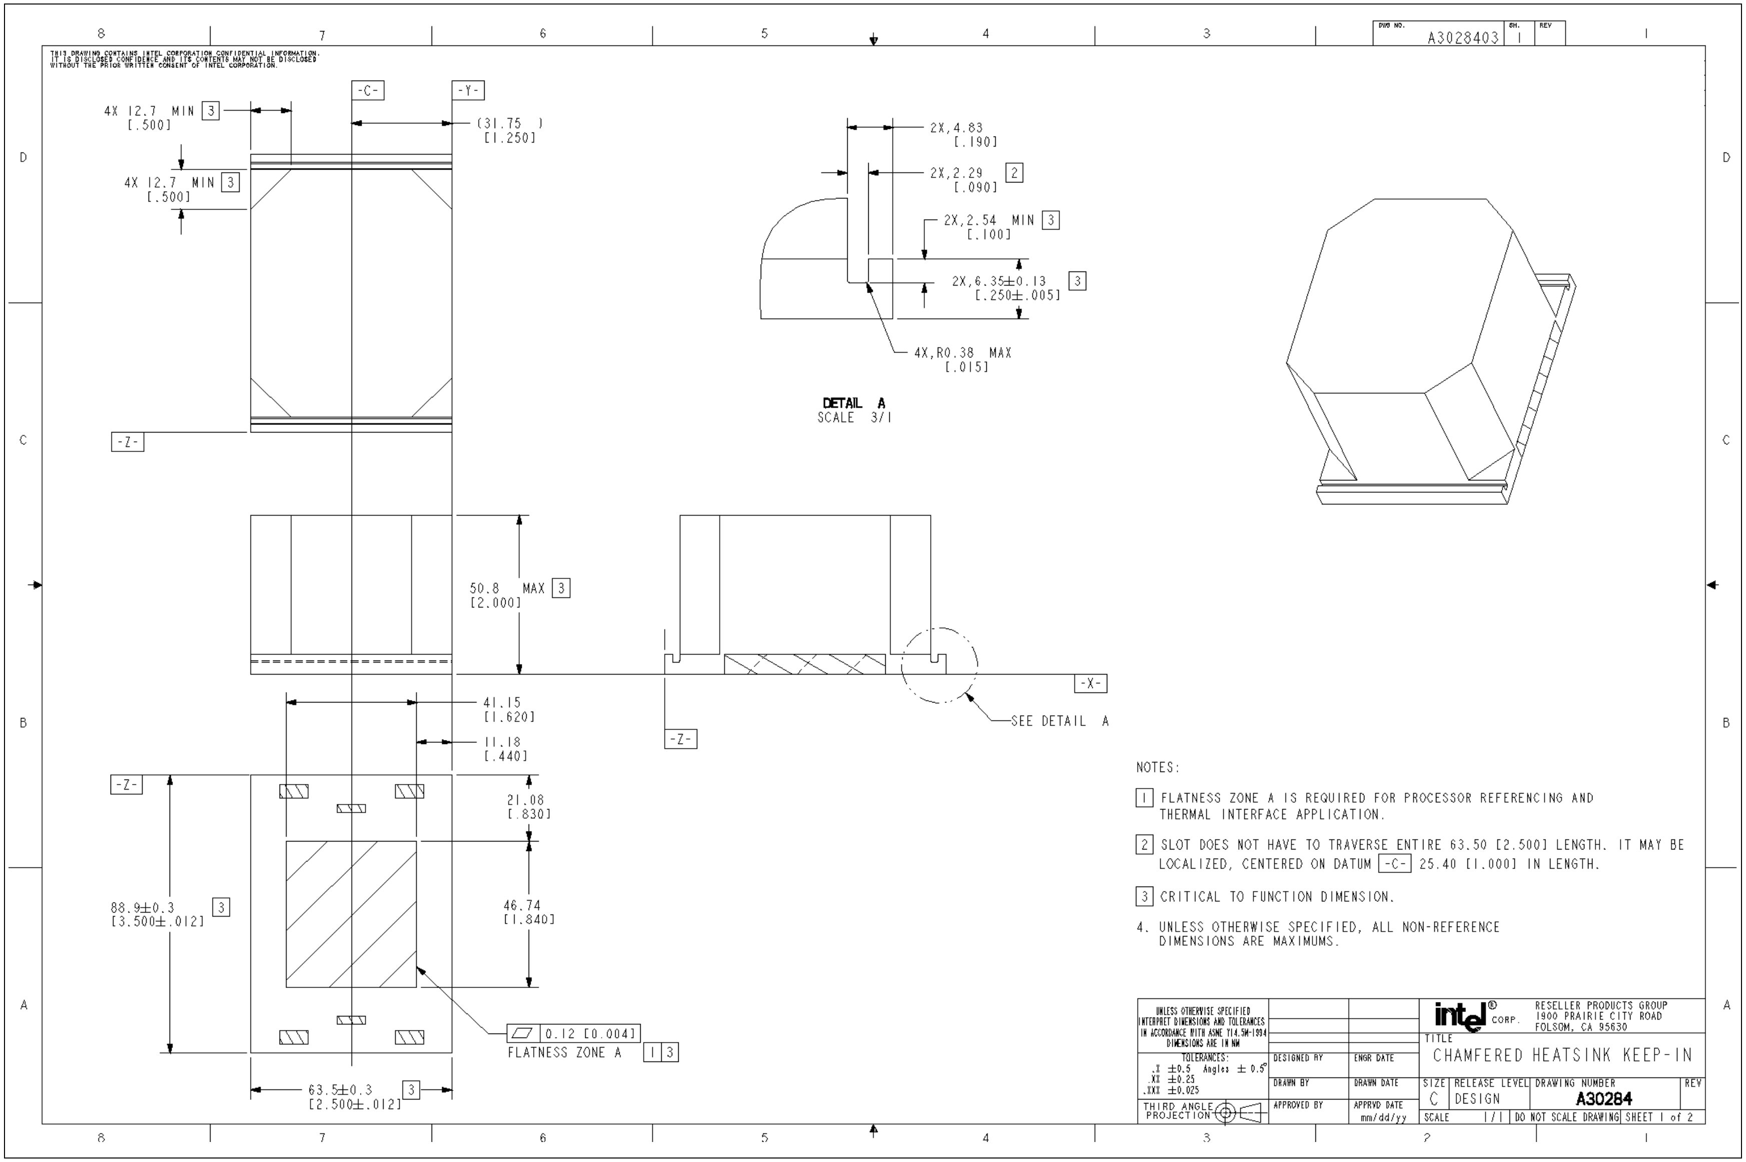

Figure C-2. Volumetric Keep-Out Zones For Airflow (Sheet 1 of 4)

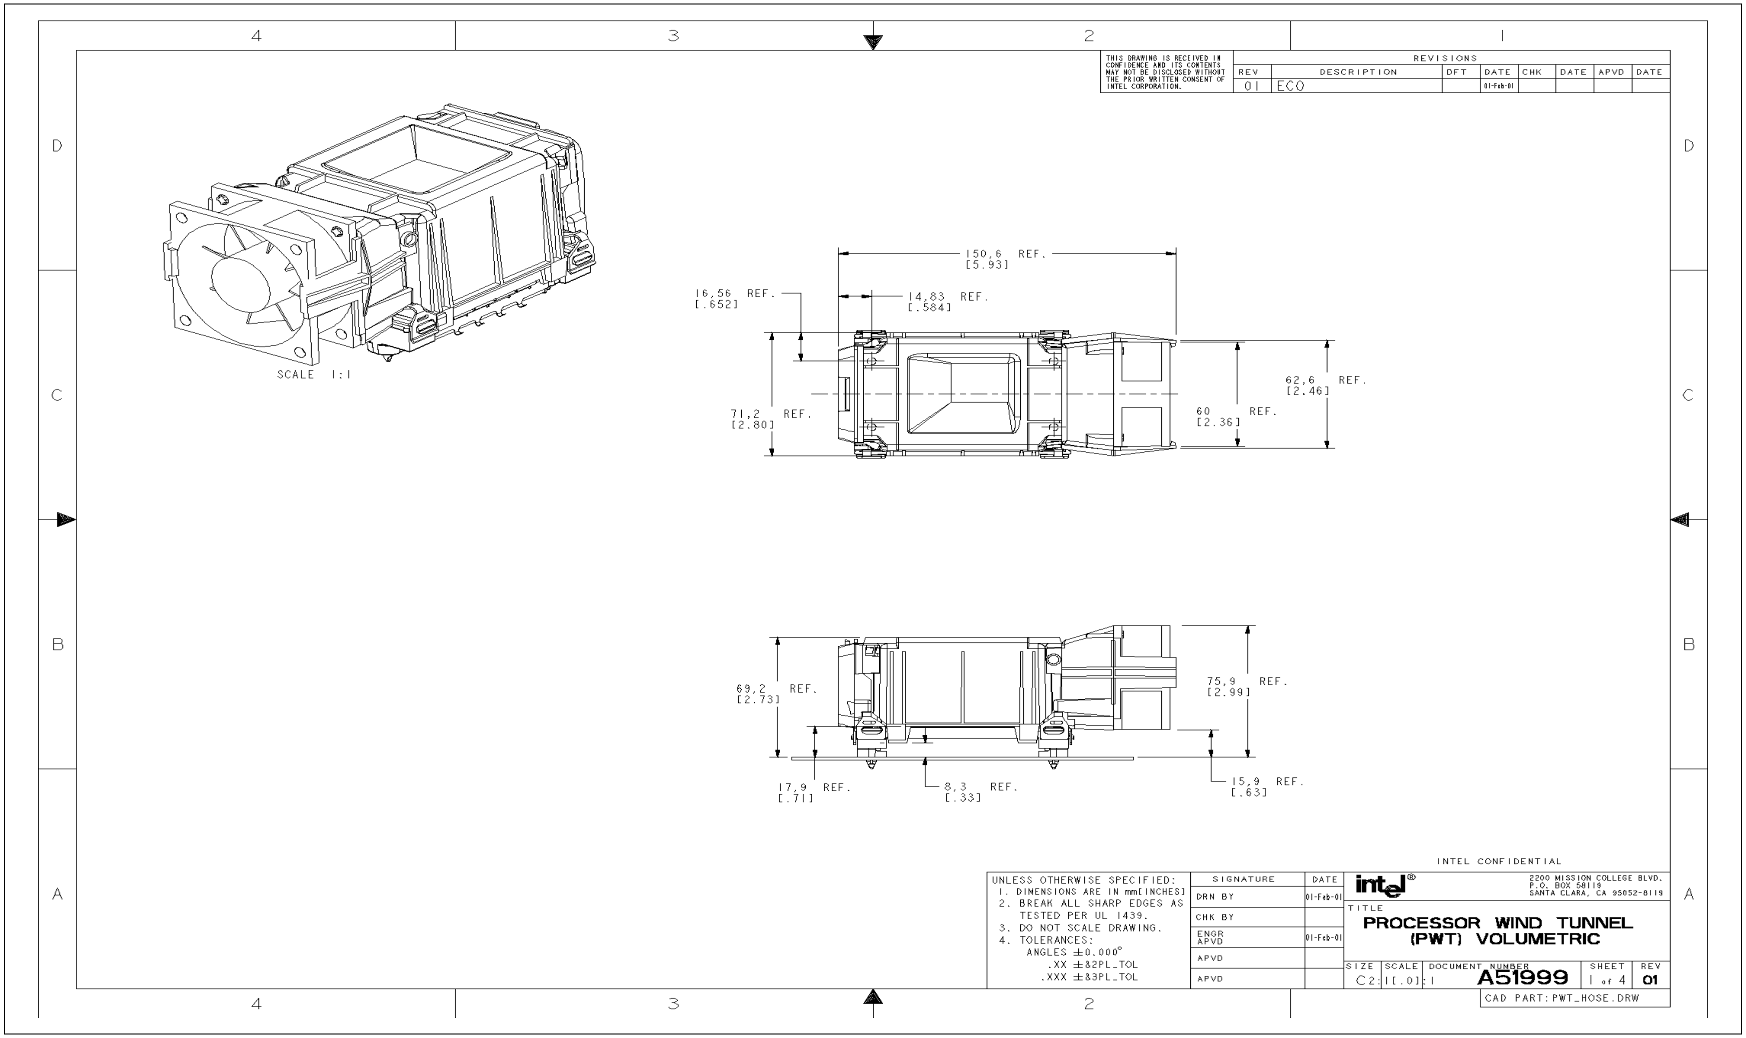

Figure C-3. Volumetric Keep-Out Zones For Airflow (Sheet 2 of 4)

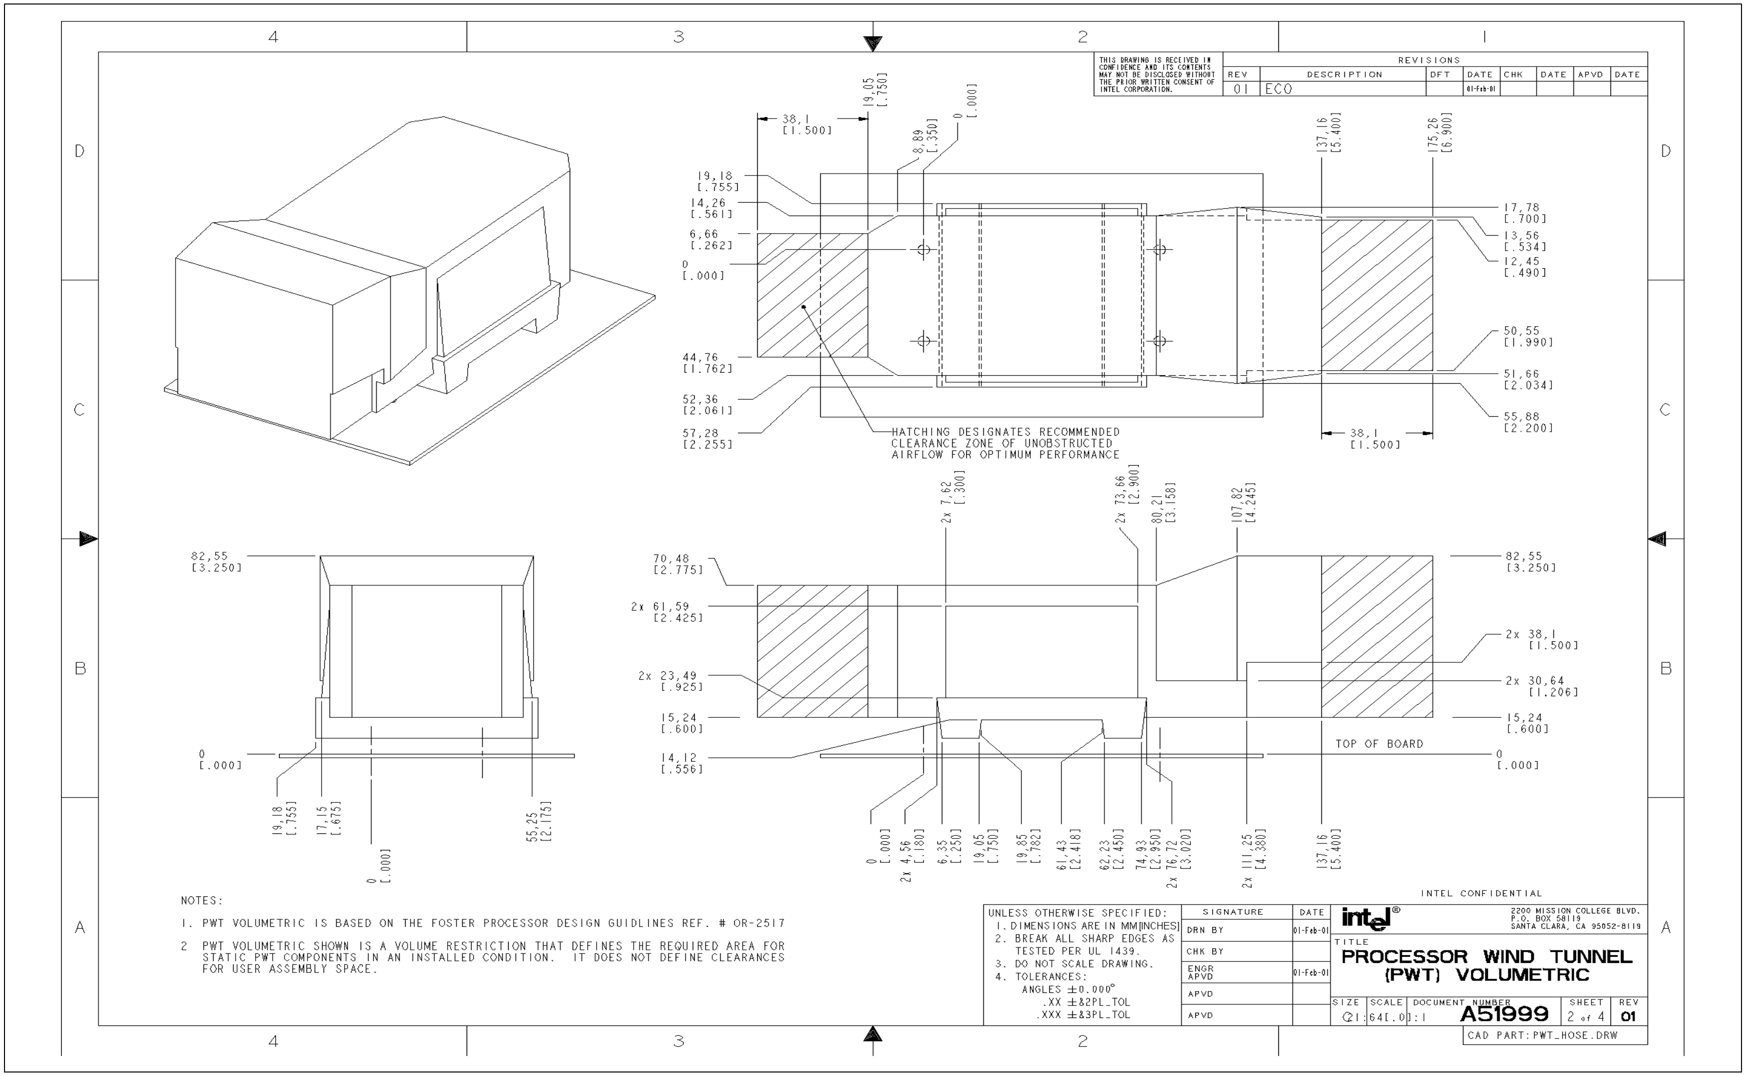

Figure C-4. Volumetric Keep-Out Zones For Airflow (Sheet 3 of 4)

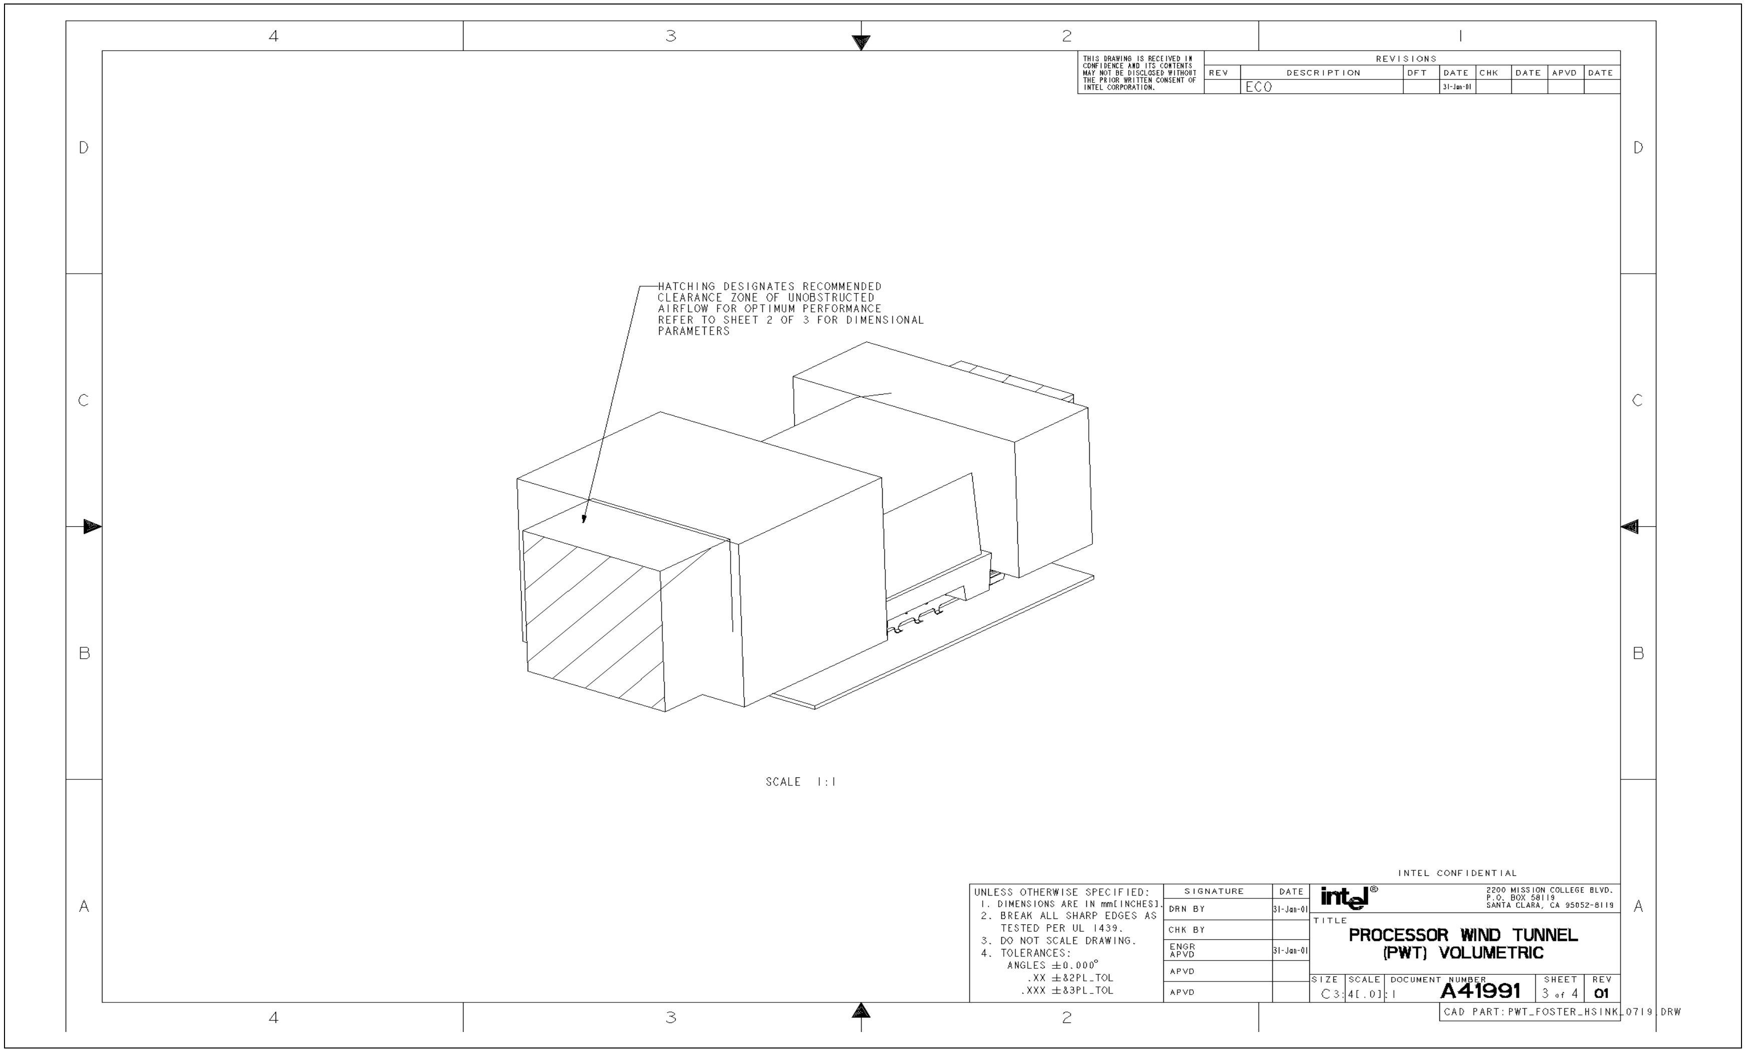

Figure C-5. Volumetric Keep-Out Zones For Airflow (Sheet 4 of 4)